

Sentiment Review Protocol

How to review any transcript, meeting, video, or document for the Ground Game sentiment tracker — granular and consistent enough that meeting #3 and meeting #30 produce genuinely comparable output. Applies to Project Ballast today; written generally enough to reuse for the next project.

1 Register the source first

Never extract from something that isn't logged. Every transcript/video/document gets a `gasplamt_sources` row (type, title, date, publisher/platform, URL if any) before anything is pulled from it. Meetings additionally get a `gasplamt_meetings` row.

2 Read the whole thing before extracting anything

One full pass, no logging mid-read. Extraction happens after you actually know what's in it — skimming-while-logging is how granularity gets lost.

3 Atomic decomposition — the mandatory first pass

The rule that caused this document to exist: when a sentence names three things together ("nitrogen oxides, sulfur dioxide, and particulate matter"), each one is its own line, not one bucketed bullet. If two different speakers cite similar-sounding studies, they're still two lines, not merged into one representative claim.

Merge/paraphrase happens only at the macro-grouping stage (step 6) — never here.

- Every distinct claim, citation, story beat, or emotional appeal gets its own line.
 - Each line is attributed: speaker name (or "unattributed" if the source doesn't make it clear), exact date, source document.
 - Direct quote or tight paraphrase — not a summary that could apply to multiple different things someone might have said.
 - Named technical/scientific/legal terms are tracked individually (this generalizes past pollutants — dollar figures, statute citations, acreage numbers, precedent names all get their own line when they're doing distinct rhetorical work).
-

4 Per-person checklist

Applied to every speaker/commenter, every time, regardless of how minor their contribution seems:

Identity — check existing stakeholders for a name match before creating a new record. If ambiguous (spelling variant, family-member confusion, ASR mis-transcription), log it as an open item in `gasplant_collection_gaps` rather than guessing. Never silently merge or fork.

Residency / distance — to the site, if stated. Flag explicitly if someone states they are NOT a local resident.

Position — oppose / support / mixed / neutral, logged in `gasplant_positions` against the specific project. This is separate from `stakeholders.stance`, which tracks Meridian's own relationship-building posture toward that person, not their public position on a given project — don't conflate the two fields.

Story — did they tell a personal story? One-line summary if yes.

Citation quality — did they cite anything, and is it real expertise/background or self-admitted casual research ("I googled it")? Note the difference explicitly, don't just log "cited a source."

Tone — assert only what the text itself evidences (word choice, delivery, pacing). Never claim a visual/audio cue (crying, gestures, tone of voice) that isn't actually in a transcript — if the source is text-only, say so rather than invent sensory detail.

Repeat appearance — did they speak more than once in this same source?

Data-center (or equivalent bait-and-switch) fear — raised or not, and the specific reasoning behind it if so — don't just tally yes/no, capture the argument shape.

Incidents — anything rising to vandalism/theft/threat/violence/harassment goes in `gasplant_incidents`, never mixed into sentiment/position data. Confidence marked verified/asserted/rumor, never overstated.

5 Officials get extra fields

- Vote, if any — logged in `gasplant_votes`, tied to the specific reading/agenda item.
 - Credibility-building moves — personal roots, site visits, tours, cited personal research — logged as their own atomic items, since these recur as a pattern across officials on both sides of a vote.
 - Direct questions asked and the answers given, especially procedural ones (what happens next, what could still change).
-

6 Only now — macro-grouping

Group the atomic list into a working set of macro categories (mid-teens is a reasonable target, adjust to what the material actually supports). Every macro category must be traceable back to specific atomic lines — no category invented from a vibe rather than the actual pile of pieces. Re-derive the grouping fresh each time rather than forcing new material into last time's categories; note explicitly when a new pattern only became visible because everything was laid out flat first (that's the point of doing it in this order).

7 Output, every time, no exceptions

- Clear version marker in the document itself (masthead/footer) and in the filename.
 - Local file: normal HTML/doc in the relevant project folder.
 - Supabase copy: matching content in `public.project_notes`, `doc_filename` following MGW.
`{category}.{kebab-slug}.{MMDDYY}.{v#}.{ext}`.
 - No Artifact tool.
 - Never delete a superseded version — move it to a `z.`-prefixed folder in the same directory instead.
 - State explicitly what's NOT yet covered (e.g. "bounded by sources collected so far, not a full crawl") rather than implying completeness that hasn't been reached.
-

Written 2026-07-15 after the Project Ballast atomic-claims and pollutant-chart exercises surfaced the need for a repeatable standard. Local copy: `MGW.protocol.sentiment-review.071526.v1.html` · Supabase: `public.project_notes`, same title. Also saved as a standing feedback memory so future sessions apply this without re-explanation.