GTM Society

powered by **Expandi**

A GTM SOCIETY PLAYBOOK

The Content System Playbook

Turn the conversations you're already having into content that doesn't sound like AI.

7 build steps | Liftable prompts | MVP in one week

By **Ferdinand Goetzen** (The Growth Syndicate) & **Glenn Miseroy** (Expandi) — from the GTM Society live session.

"AI raised the floor and lowered the ceiling. Your transcripts are how you get the ceiling back."

THE REFRAME

AI raised the floor. It also lowered the ceiling.

Everyone can generate content at scale now — that was never the problem. The problem is the output: samey, "good enough," never unique. And readers can feel it. The 2000s gave people banner blindness, then email blindness. Now they have **AI blindness** — they scroll past slop without registering it.

Ferdinand once posted an obvious joke about "pitching TGS to the king" at a Madrid conference — and got a comment earnestly praising his bravery. When the only engagement comes from people who clearly didn't read, that's what slop earns you.

Here's the part almost every team misses: your people are having genuinely nuanced conversations all day. Real opinions, real points of view, real frameworks — and thanks to note-takers, most of it is **already documented**. Then they open Claude and generate content from nothing.

The insights that would make your content exceptional already exist in your company. This playbook is the system for connecting them to your output.

> **"You're sitting on a goldmine of unique, differentiating points of view. You just need to create the connection."**

PROOF

What this looks like when it works.

All numbers from The Growth Syndicate running this exact system.

0 → ranking
IN UNDER 2 MONTHS

TGS did **no SEO at all**, then switched on SEO + AEO with this system — results came in weeks, not quarters.

3 viral posts
IN ONE MONTH

Three LinkedIn-viral posts against a normal average reach of 2–3K — **all generated through the system**.

6–8 wks → ~1
REPORT WRITING TIME

The writing of a data-heavy industry report dropped from the bulk of a 6–8–week project to about one week. Gathering the data and interviews still takes its time — **the system collapses the writing, not the research**.

BOTTOM LINE

One person runs the content output of what used to take a **~10–person team**.

They built the sophisticated version over months. You can build the MVP **this week** — that's what the next pages are.

HOW TO USE THIS

A system, not a stack of tips.

WHO IT'S FOR

GTM engineers
who can wire the tools but want the output to sound human

Founders
doing their own content

Marketers
who want AI speed without AI smell

THE SYSTEM, IN ONE LINE

Centralize inputs → Keep them fresh → Parse into tagged insights → Lock the voice → Generate with injection → Govern what goes in

HONEST EXPECTATIONS, UP FRONT

5–10%
quality traded away vs. your best writer, hand-writing

3–8×
speed gained on everything you ship

2 hrs
to a 6/10 system. A 7/10 takes a week or two.

Getting from 7 to 9 takes months — that's the part you **don't need on day one**.

01 STEP ONE

Centralize your inputs in one knowledge base

WHY IT MATTERS

The system is only as good as what flows into it, and input scattered across five tools flows nowhere. You need one place where everything sinks in automatically — **one memory**. Notion is the recommendation: not because it's fancy, but because it integrates with practically everything (note-takers, Slack, automations) and your generation tool can read it.

DO THIS

1. Create one Notion base as the knowledge hub. This is the **memory, not the brain** — Claude is the brain.
2. Connect your note-taker (Fireflies or equivalent) so meeting transcripts land there automatically.
3. Add a drop zone for voice notes and one for ideas (a piped-in Slack or Discord channel where the team already debates things counts).
4. **Handpick what goes in at the start.** Do not connect everything you have.

MEMORY VS. BRAIN

Keep the roles straight: the base **stores and organizes** what your company knows; the model **reasons and writes** from it. Everything in this playbook flows from that separation — you improve output by improving what's in the memory, not by prompting harder.

01

STEP ONE · CONTINUED

Centralize your inputs in one knowledge base

SWIPE THIS THE STARTER BASE — COPY

Notion table, five columns:

Source (call / voice note / idea channel / creator post)
Date
Contributor
Raw content or transcript link
Status (unparsed / parsed / excluded)

That's the whole schema. The value is in what you put through it, not the template.

WATCH OUT — CONTEXT OVERLOAD IS REAL

More input is not better. Past a point, noise makes the model's answers **worse**, not richer — and the higher your company's volume (leads, users, meetings, content), the faster you clog it. Start with a curated trickle, not a firehose.

02 STEP TWO
Create input rituals — where solo founders win

— WHY IT MATTERS

A knowledge base with stale content produces stale output. The best single source, run by the people who built the system: a **recurring founder/stakeholder interview** — one session where someone with opinions talks through what they're seeing. Input doesn't have to be structured. It has to be *fresh*.

— DO THIS

1. **Monthly:** a 30–60 min interview with each founder/key stakeholder. One person asks, the note-taker documents. This will be your richest vein.
2. **Weekly:** voice notes. Brain-dump observations, hot takes, things a prospect said that surprised you. Whisperflow or any voice-to-text works.
3. **Solo founder with no meetings?** Voice notes ARE your input engine. Daily musings are enough — the system turns rambling into structure; that's its job.

SWIPE THIS WEEKLY VOICE-NOTE PROMPT COPY

```
"What did I see this week that most people in my market get wrong? What did a
customer/prospect say that surprised me? What am I convinced of that I wasn't 3 months
ago?"
// put it in your reminders
```

03 STEP THREE

Borrow taste while you build your own

WHY IT MATTERS

Early on, your own insight base is thin. The bootstrap: pick a handful of creators whose worldview matches yours and feed their public thinking in as *reference* — emulate, not copy. TGS started exactly this way — a half-dozen well-known GTM and positioning creators, ~50 posts each — and then **removed them** once their own original ideas filled the base. The training wheels are meant to come off.

DO THIS

1. Pick 3–5 LinkedIn creators genuinely aligned with your worldview — people you'd want to sound like **in spirit, not in phrasing**.
2. Scrape ~50 posts each (PhantomBuster, Expandi's post tools, or any post scraper) on a recurring basis.
3. Load them into the knowledge base tagged as *external reference*, separate from your own insights.
4. Review quarterly: as your own insight count grows, cut the external references down. The goal is **your voice, informed — not theirs, laundered**.

WATCH OUT

If your output starts reading like a specific creator's feed, you've crossed from emulation to copying. The tag separation in step 3 exists so you can **dial external influence down**.

04 STEP FOUR

Parse inputs into tagged insights — the step everyone skips

WHY IT MATTERS

Raw transcripts are noise. The system works because every input gets cut into **discrete insights, each tagged** — and this is also the moment your content strategy stops being a slide nobody opens. Your content pillars become the actual tags on actual insights. AI, ironically, makes you *more* aligned to strategy, because the strategy is operationally wired into the pipeline.

DO THIS

1. Define your insight types. TGS uses: **insight · framework · trend observation · methodology.**
2. Define your core topics — the 3–5 themes you want to own (theirs: talent & culture, strategy & positioning, measurement & ROI). Add secondary tags as subtopics.
3. Run every new transcript/voice note through a parsing prompt (next page). Store the output rows, not the raw transcript.
4. Tag every insight with its **contributor**. Attribution matters later — it's what lets you auto-inject "real quote from a real person" into content.

WHAT A PARSED ROW LOOKS LIKE

Insight: "Most teams over-invest in creation and under-invest in distribution." · **Quote:** verbatim from the call · **Contributor:** Ferdinand · **Type:** trend observation · **Topic:** strategy & positioning · **Subtopic:** content ops

04 STEP FOUR · CONTINUED
Parse inputs into tagged insights

SWIPE THIS THE PARSING PROMPT COPY

```
Here is a transcript. Extract every distinct, concrete insight — an opinion, framework, observation, or method, in the speaker's own words where possible.

For each, return:
· the insight (2-3 sentences)
· a verbatim quote
· contributor
· type (insight / framework / trend observation / methodology)
· core topic (from this list: [your pillars])
· subtopic

Skip small talk, logistics, and anything generic enough that a competitor could have said it.
```

THE LINE THAT DOES THE WORK

"Anything generic enough that a competitor could have said it" is the filter that keeps slop out of the base. If an insight would survive on a rival's blog unchanged, **it isn't yours and it isn't differentiating** — don't store it.

05 STEP FIVE
Lock the voice — the anti-banality layer

— WHY IT MATTERS

This is the single biggest win, and the one GTM engineers skip because it isn't an automation problem. Model output defaults to AI-isms — "it's not just X, it's Y," em-dash confetti, the same twenty words everyone now flinches at. A written rules layer (TGS builds theirs as JSON files) strips that out at generation time. Theirs took ~6 months of tweaking; you get the compounding benefit from day one of using *any* serious version, because it works in any tool — literally drop it into a Claude or ChatGPT chat alongside your request.

— DO THIS

1. Write (or adopt) a voice rules file: banned phrases and patterns, sentence-length rules, how strong opinions get stated, what your brand never says.
2. Include it in **every** generation call — chat, API, or Notion AI. No exceptions; consistency is the mechanism.
3. Treat it as living: every time output smells like AI, the fix goes into **the file**, not just the draft.

THE SHORTCUT — FREE FOR MEMBERS

The Growth Syndicate's **Anti-Banality Framework** — the actual JSON rules layer behind the results on page 3, in a neutral fill-in-the-blanks version — is free for GTM Society members. That's the gated bonus for this session, and members keep getting the updated versions as TGS improves it.

06 STEP SIX

Generate with injection — and keep a human steering

WHY IT MATTERS

Generic AI listicles all look identical: intro, tool #1 (always the author), tool #2, done. What makes system-generated content read as *sophisticated* is **injection** — real verbatim quotes from your contributors, original statistics from your own data, internal links placed by the system. A piece that quotes a real founder and cites a number nobody else has doesn't pattern-match to slop. Original stats are also disproportionately powerful for AEO — answer engines love a number with a source.

DO THIS

1. Generate from the insight base, never from a bare prompt. The ask is **"write X using these tagged insights,"** not "write X."
2. Inject at least one real quote (with attribution) and, where you have it, one original stat per piece. If you're a SaaS company, anonymized product data is an untapped stat mine.
3. Keep human steering: the system drafts, a person directs and makes the final pass. The 15-minute publish is real — but it was 15 minutes of *steering*, not zero minutes.
4. Reuse the base across formats: the same insight cluster becomes an SEO piece, a LinkedIn post, a report section, a landing page.

WATCH OUT — THE QUALITY TRADE IS HONEST

Roughly 5–10% below what your best writer produces by hand, at 3–8× the speed. If a piece matters enough that 10% matters, **write it by hand**. The system exists so that everything else still gets made.

07 STEP SEVEN
Govern what goes in

WHY IT MATTERS

The failure mode nobody plans for: sensitive material leaking into public content. Connect every note-taker to the pipeline and eventually a C-level call, a personnel discussion, or internal strategy ends up **one careless generation away from a LinkedIn post**. The fix is boring and non-optional: an explicit rule for what enters the system.

DO THIS

1. **Default-exclude:** internal strategy, personnel, legal, and C-level calls stay out unless explicitly cleared.
2. Make inclusion a decision: a weekly 5-minute pass — *"which of last week's meetings go to the system?"* — beats auto-ingesting everything.
3. Anonymize data inputs (product stats, client examples) before they enter the base.
4. Put **one person's name** on this. Moderation without an owner is moderation that doesn't happen.

THE PLAN

The MVP: your first week.

Day 1 2 HOURS

Notion base up. Note-taker connected. Voice rules file in place (members: grab the Anti-Banality Framework on page 15). **This is the 6/10 version — and it already beats generating from nothing.**

Days 2–3

Pick your 3–5 creators, scrape ~50 posts each, load as external reference. Record your first voice note. Define insight types + core topics.

Days 4–5

Parse your first 3–5 transcripts with the parsing prompt. You now have a real insight base.

End of week

Generate your first piece *from the base*, with one real quote injected. Compare it to what you were shipping before. That difference is the whole argument.

CALIBRATION

A 7/10 system takes another week or two of tweaking. 7 → 9 took the pros months — **you don't need it to start winning.**

APPENDIX

The swipe file.

01 THE STARTER BASE COPY

Notion table, five columns: Source · Date · Contributor · Raw content / transcript link · Status (unparsed / parsed / excluded)

02 WEEKLY VOICE-NOTE PROMPT COPY

"What did I see this week that most people in my market get wrong? What did a customer/prospect say that surprised me? What am I convinced of that I wasn't 3 months ago?"

03 THE PARSING PROMPT COPY

Extract every distinct, concrete insight. For each: the insight (2-3 sentences) · verbatim quote · contributor · type · core topic ([your pillars]) · subtopic. Skip anything generic enough that a competitor could have said it.

THE GATED BONUS

Grab the Anti-Banality Framework — free for members.

The actual JSON voice layer behind the results in this playbook, in a fill-in-the-blanks version — and members get every improved version as it evolves. Join free, grab it, and post your first system-generated piece in the community.

Join GTM Society →

Want the full walkthrough? **The session recording — Ferdinand walking through the real system, screens and all — lives in the community.**

See you in there.
Ferdinand & Glenn · GTM Society — powered by Expandi