

## Meeting-to-Content

### Inputs

Workflow — Text (blog + 2 LinkedIn posts)

1. Fetch the meeting
2. Find the upshot
3. Pull the strongest moments
4. Draft the three text assets
5. Voice + grounding pass
6. Save the text files

Workflow — Clips (3–4 short videos)

1. Pick the clip moments
2. Get exact timestamps
3. Create the clips in Grain
4. Wait for approval
5. Bundle into a story for download
6. Report back

Voice guide

Format specs

Blog post (2,000–2,800 words)

LinkedIn posts

Failure modes

# Meeting-to-Content

Turn a Grain recording into:

- **1 blog post** (2,000–2,800 words)
- **2 LinkedIn posts**
- **3–4 short video clips** (30–45 seconds each)

Text assets are saved as finished markdown files. Clips are created in Grain, approved by the user, then bundled into a story for download.

## Inputs

- **Grain meeting identifier (required):** a URL (e.g. <https://grain.com/share/recording/8421e934-...>) or a meeting ID.
  - **Audience + goal (required):** one line naming who the content is for and what you want them to do — e.g. *“For people new to Grain — get them to try the app.”* If it’s missing, ask before drafting. Don’t guess.
  - **Angle (optional):** one line on which moment or idea to lead with. If omitted, derive it from the meeting.
- 

## Workflow — Text (blog + 2 LinkedIn posts)

### 1. Fetch the meeting

Use the Grain MCP:

- `fetch_meeting` — title, attendees, date, duration
- `fetch_meeting_transcript` — full transcript with timestamps
- `fetch_meeting_notes` — if notes exist

If the meeting isn’t accessible, stop and tell the user.

### 2. Find the upshot

Read the transcript and write one plain-English sentence for what the meeting actually said. A smart friend who didn't attend should get it. It should reference something specific (a workflow, a quote, a surprising point). No jargon unless it's quoted from the meeting. This sentence anchors everything.

### 3. Pull the strongest moments

Scan the transcript for opinionated statements, concrete stories with specifics, surprising points, and memorable verbatim quotes. For each, note the speaker, the timestamp, the quote, and why it's strong. Aim for 8-10 moments — you'll reuse this list for the clips too.

Assign them so the assets don't repeat:

- **Blog post** synthesizes across all moments.
- **LinkedIn post #1** leads with one strong moment.
- **LinkedIn post #2** leads with a *different* moment and a different shape than #1.

### 4. Draft the three text assets

Follow the **Format specs** below.

- **Blog post** (2,000–2,800 words): the canonical version. Include a TL;DR near the top and an FAQ (8–10 questions) near the bottom — these are what make it citable by AI search. Carry the CTA to the goal from your inputs.
- **LinkedIn post #1 and #2**: short, direct, one idea each. Each stands on its own.

### 5. Voice + grounding pass

Rewrite every draft against the **Voice guide** below, then adapt by format (blog = more depth; LinkedIn = short and punchy). Then check:

- Every claim, quote, and number exists in the transcript. If a line can't be tied to a moment, cut it or fix it.
- Plain language. No corporate filler.
- **No links in LinkedIn post bodies.** If a post needs a link, note it separately as a first comment.
- **LinkedIn posts don't name people.** Attribute quotes by role ("our CEO," "an engineer on the team").

### 6. Save the text files

Write three markdown files to the output folder: blog-post.md, linkedin-1.md, linkedin-2.md.

---

## Workflow — Clips (3-4 short videos)

### 1. Pick the clip moments

From the moments list in step 3 above, choose the **3-4 that work best as standalone 30-45 second clips**. A good clip moment:

- Is a single continuous stretch of one speaker talking (no need to stitch cuts)
- Runs roughly 30-45 seconds start to finish
- Makes sense with no setup — it opens and closes on complete

thoughts

- Is opinionated, surprising, or a concrete how-to worth watching on its own

Pick different speakers or different ideas across the clips so they don't feel repetitive.

## 2. Get exact timestamps

For each clip, find the **start and end timestamps** in the transcript. Anchor on the verbatim quote: locate the exact opening line and closing line in the transcript, confirm they appear word-for-word, and read the timestamps off those lines. The clip must start and end on complete thoughts. Add a small buffer (a second or two) at each end so nothing is clipped.

## 3. Create the clips in Grain

For each clip, call `create_clip` with the recording ID, start time, end time, and a short descriptive title (the hook of the clip). Create all 3-4. Each response returns a `clip_url` — record it; it's required for the next step.

## 4. Wait for approval

**Stop and show the user the clips you created** — title, timestamp range, the one-line hook, and **the `clip_url` returned by `create_clip`** for each. The link is not optional: the user can't preview or share a clip without it, so if a clip table is missing links, the step isn't done. Ask the user to confirm which ones to keep. Don't proceed to the story until they approve. If they want changes, adjust the timestamps and recreate.

## 5. Bundle into a story for download

Once approved:

1. Call `create_story` with a title (use the meeting title + "clips" or the upshot).
2. Call `add_clips_to_story` with the approved clip IDs.
3. Give the user the story link so they can download the clips.

## 6. Report back

Short summary: the source meeting, the one-sentence upshot, the three text file paths, and the story link with the clip count.

---

# Voice guide

Apply this to every text asset before format-specific tweaks. The voice is a real operator talking to a peer, not a vendor talking to a buyer.

- **Plain language.** If you wouldn't say it in a Slack DM to a coworker, rewrite it.
- **Specific over general.** Names, numbers, real examples, real quotes. Not "leverage," "synergy," or "solutions."
- **One idea per sentence.** Short paragraphs. Let sentence length vary so it doesn't read like a machine.
- **Active voice, present tense.**
- **No em dashes.** Use a period or a comma instead. This is a hard rule.

- **No corporate filler.** Cut “I’m excited to share,” “thrilled to announce,” “in today’s fast-paced world,” “in our recent webinar,” and anything like them.
  - **Declarative, not hedgy.** Say the thing. Drop “I think,” “it seems,” “arguably” unless the uncertainty is the point.
  - **Avoid these words unless you’re quoting the meeting:** agentic, operating model, transformation, compounding, ecosystem, paradigm, leverage (as a verb), unlock (as a noun), second brain, delve, robust, seamless.
  - **Use the customer’s vocabulary.** For a general or new-user audience, say “connector” not “MCP,” and “knowledge base” not “knowledge graph.” For a power-user audience, the technical terms are fine.
- 

## Format specs

### Blog post (2,000-2,800 words)

Structure, top to bottom:

1. **Title** — clear and specific, no clickbait.
2. **TL;DR** — 2-4 sentences at the very top summarizing the whole post. This is what AI search engines quote, so make it self-contained.
3. **Intro** — the upshot, framed for the audience and goal, in 2-3 short paragraphs.
4. **Body** — the meeting’s ideas organized into 3-5 sections with descriptive headers. Ground every section in a specific moment, quote, or example from the transcript. Where you introduce a concept, define it plainly in one sentence (“What it is: ...”).
5. **FAQ** — 8-10 real questions a reader would ask, each with a short direct answer. This section, with the TL;DR, is what makes the post citable by AI search.
6. **CTA** — one clear next step tied to the goal (e.g. try the app). The blog is the only asset that carries an outbound link.

### LinkedIn posts

- **Length:** short. Roughly 100-200 words. Front-load the hook in the first line.
  - **Shape:** one idea, told as a short story or a clear argument. No headers, no bullet lists unless it genuinely helps.
  - **No link in the body.** If there’s a link, it goes as a first comment.
  - **No names.** Attribute quotes by role or relationship.
  - **CTA:** soft and specific, or none. Don’t hard-sell.
  - Post #1 and #2 lead with different moments and take different shapes (e.g. one narrative, one contrarian take).
- 

## Failure modes

- **Meeting too short or low signal:** say so. Don’t force content out of a 15-minute internal sync.
- **Grain MCP unavailable or not authorized:** stop. The skill needs the transcript and clip tools.
- **Transcript is mostly filler:** use the strongest 20-30% and treat the rest as scaffolding.
- **No moment cleanly fits 30-45 seconds:** it’s fine to ship 3 clips instead of 4. Don’t force a clip that needs setup or trails off mid-thought.

