MY WEIRD PROMPTS · PIPELINE CHANGELOG

Q4 2026

October – December 2026 · 8 changes

Quarter in progress — covers changes to 14 November 2026.

Every episode of My Weird Prompts is written, voiced and assembled by an automated pipeline. This changelog records its major updates — new models, new voices, new stages — and the first episode each one produced. Small prompt tweaks happen weekly and are not listed.

8 October 2026 · Infrastructure · Pipeline V5.2 · Prompts 4.1

Episode Books, an MCP server and an AI sandbox

Three ways into the archive that don't involve listening. Every episode now has an Episode Book: a PDF with the episode's record, synopsis, transcript and sources, pre-rendered for the whole archive. AI tools can query the show directly through a public, read-only MCP server at myweirdprompts.com/mcp. And /sandbox/ hands out a condensed version of the show's format prompt and character sheets, so anyone can run their own version of the show in a chatbot.

FIRST EPISODE AFTER RELEASE: #5775 What "Leaked System Prompts" Actually Mean

8 October 2026 · Writing · Pipeline V5.2 · Prompts 4.1

Hilbert gets ridiculous, and deep topics get thirty minutes

Hilbert's stories are now meant to be absurd, landed through exact detail and a straight face, where before they had to be believable. His turn grew to three or four minutes. Episodes on topics that need it can now run to thirty minutes. The planner had been offering that length all along, but the code quietly cut it back to twenty-five.

- Later the same day (prompts 4.3), Hilbert started treating Corn and Herman as colleagues he knows. His crankiness is aimed at the world, not at them.

FIRST EPISODE AFTER RELEASE: #5776 Why UV Epoxy Keeps Curing in the Dark

8 October 2026 · Infrastructure · Pipeline V5.2 · Prompts 4.2

Anyone can send in a prompt

A public Telegram bot, t.me/mwplistenerbot, takes prompts from listeners. It asks four questions (your name, a line about you, where you are, and the question itself) and emails the submission to me. Nothing becomes an episode automatically: I read each one and decide. Episodes now close by reading out the bot's address instead of an email address.

FIRST EPISODE: #5783 Can a Dremel Be Gym Equipment?

9 October 2026 · Voices · Pipeline V5.2 · Prompts 4.6

The short-code slate is retired

Hilbert no longer reads out the short code and title at the start of each episode. The slate cost an extra GPU step on every episode, and I decided it wasn't worth it. The /e/ short links it advertised still work.

FIRST EPISODE AFTER RELEASE: #5796 900 Tonnes on Public Roads: The Weirdest Cargo Moves

9 October 2026 · Writing · Pipeline V5.2 · Prompts 4.6

The sign-off suggests related episodes

Just before the sign-off, the hosts now name up to three earlier episodes on related topics, by number and title. The suggestions come from a similarity search over the archive and are inserted by code, so the numbers and titles are always real ones. The script model is told not to invent its own.

FIRST EPISODE: #5809 What If Israel Had Never Fallen? A 2,000-Year Thought Experiment

10 October 2026 · Writing · Pipeline V5.3 · Prompts 4.7

V5.3: better recommendations

Episodes that were being recommended almost every time are now dropped from the related-episode suggestions. One episode had turned up in 12 of 19 recent runs. The recommendation is also placed directly before "This has been My Weird Prompts" rather than at the start of the host's line. And the script writer now treats a word in my prompt that doesn't fit as a likely speech-to-text error, rather than building a claim on it.

- New episode addresses now end in a random five-digit number, so two episodes on the same topic can't end up with the same URL. Existing addresses are unchanged.
- Later the same day the site gained an entity index at /indices/: the people, organisations, places, works, events and products the show discusses, each linked to its episodes.
- Also later that day, the voices briefly moved to cheaper T4 GPUs, then back to L4 once the speech decoder was rebuilt around a CUDA graph with an fp16 backbone. On the L4 it renders nearly three times faster than before, for about $0.28 of GPU time per hour of audio, down from $0.77. It is on trial: if the voices drift, it can be switched back.

FIRST EPISODE: #5835 What Do You Call the Urge to Know Obscure Things?

24 October 2026 · Script model · Pipeline V5.21 · Prompts 4.29

V5.21: every model brought up to date

I audited every model the pipeline calls apart from the script writer and brought them all up to date. Where a stage only needs a good fast model, it now asks OpenRouter for the latest version of its family, so it moves to a new generation without anyone touching the code. Titles, descriptions and the blog post come from Gemini Flash, as structured output with length limits, and the blog post is much shorter. Covers are drawn by a newer image model that is shown earlier covers as a style guide.

- Title, description, URL and image prompt: one Gemini Flash call that returns fields rather than free text, with limits on each (title up to 60 characters, description 50 to 90 words). A miss is regenerated. Before this, one DeepSeek call wrote all of it plus the blog post in labelled text.
- Blog post: a separate call, about 100 words instead of 150 to 220.
- Voice-memo transcription: Gemini Flash, latest version, instead of a Gemini 3.1 Flash Lite preview. Whisper is still the fallback.
- Search embeddings: the released Gemini Embedding 2 instead of its preview. The vectors are identical, so nothing had to be re-embedded.
- Covers: Krea 2 Turbo Style, given three earlier covers as style references, about a cent per cover. Flux Schnell, the previous model, is the fallback.
- Embeddings stay on a fixed version, because a silent change there would invalidate every stored vector.
- The experiment and interview episode types were removed: the guest models they relied on had been retired.

FIRST EPISODE: #6245 Why Inference Engines Change Benchmark Scores

24 October 2026 · Script model · Pipeline V5.22 · Prompts 4.29

V5.22: the script writer follows GPT Sol's latest version

The script writer now asks OpenRouter for the newest GPT Sol rather than naming GPT-6.1 Sol directly. Today that is still 6.1, so nothing changes yet, but when OpenAI ships a new Sol the show moves to it without a code change. The credit at the end of each episode still names the exact version that wrote the script.

FIRST EPISODE: #6245 Why Inference Engines Change Benchmark Scores