Episode Book

EPISODE #6637

EQ, Compression, and Limiting for Spoken Voice

PUBLISHED	16 November 2026 at 21:27 UTC
DURATION	21:28
LISTEN	myweirdprompts.com/episode/eq-compression-limiting-voice-09826/
HOSTS	Corn, Herman, Larry, Hilbert
SCRIPT MODEL	z-ai/glm-5.3
VOICES	omnivoice
PIPELINE	V5.34
SOURCES	20 cited by the research (listed at the end)
BOOK GENERATED	16 November 2026 at 22:22 UTC

SYNOPSIS

A practical refresher on the three tools that do almost all the work in spoken-voice editing: EQ, compression, and limiting. The chain order matters—cutting rumble below eighty hertz and sweeping out low-mid mud before compression keeps background noise from triggering the compressor. Rather than obsessing over compressor knobs, watching for three to six decibels of gain reduction on the loudest words provides a reliable target. The episode covers protective brickwall limiting, REAPER's stock plugins, and why Apple Podcasts' minus sixteen LUFS target applies to both mono and stereo.

TRANSCRIPT



Corn

Daniel's opening line got me. He says he sat down for a couple of editing lessons with a friend who's a real music producer, and spent the whole time "staring at these fancy knobs" — attack, gain, release — with no idea what any of them did. This was years ago, for voice-over work on his YouTube videos. The friend's DAW of choice is REAPER, which suited Daniel fine because he's on Linux. And somewhere in those sessions he discovered something that I think surprised him: that dragging points along an EQ curve is, in his words, weirdly addictive.



Herman

I'm Herman Poppleberry, and I want to say immediately that the addictive part is not a side note. It might be the most honest thing in the whole prompt.



Corn

He's asking for a refresher. The knowledge has faded, he says, and he wants a run through of the basic tools of digital audio editing — the commonly reached-for levers, especially for voice. So that's the episode: the three tools that do almost all the work, the handful of numbers worth memorizing, and where to stop.



Herman

Three tools, and the easiest way to hold onto them is what each one changes. EQ changes tone — the balance of low, mid, and high energy in the voice. Compression changes how steady the level is, how far the loud words sit above the quiet ones. And the limiter changes nothing at all, on purpose. It's a ceiling.



Corn

A ceiling with a security guard standing under it.



Herman

There's a line I love from a journalism audio-training course: used lightly together, EQ and compression turn a decent recording into something that sounds like radio. Used heavily, they turn it into a fast-food drive-through speaker.



Corn

Which everyone has heard, and nobody has enjoyed.



Herman

Right. And the through-line of everything today is that each of these tools is a small number of controls with a physical meaning. Threshold is a line the signal crosses. Ratio is how hard you push back once it's over. Attack and release are timing. None of it is magic, all of it is movable — and there are sensible starting numbers for spoken voice, so you're not guessing in the dark.



Corn

The REAPER part matters here too. It runs natively on Linux, it ships with every plugin you need for this, and — this is the part I like — it has no opinionated default channel strip. There's no preset mojo baked in. Every bit of polish is a decision you made.



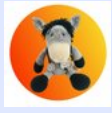
Herman

Which sounds like a burden and is actually a mercy.



Corn

Because every problem is one you can find and remove. Nothing is hidden under factory color.



Herman

So, EQ first. And not just because it's first in the alphabet — there's a mechanical reason the order is EQ into compression rather than the other way around. If you compress a voice that's still full of rumble, the compressor reacts to the rumble. It's sitting there measuring everything below a hundred hertz and squeezing your consonants because a truck went past. Clean the tone first, then even out the level.



Corn

What does the chain actually look like?



Herman

High-pass at eighty hertz is the anchor point. Seventy for a deep male voice, up to a hundred for a light one. Most speaking voices carry almost nothing useful down there, and that's where rumble, desk thumps, and the heating system live. Then the classic problem zone: two hundred to four hundred hertz, the low mids. Home recordings pile energy up there and the voice starts sounding thick, cloudy —



Corn

Talking into a cereal box.



Herman

That's the exact phrase the training course uses, so take it as official. Then presence at three to five kilohertz, where the consonants live and where the sense of someone talking to you rather than near you comes from. Then sibilance at five to nine kilohertz, handled after compression, and optionally a touch of air above ten.



Corn

The cereal box is going to haunt me. How do you find the mud, though, if you don't know where it is in your voice?



Herman

The sweep. This is the one technique I'd tattoo on Daniel's forearm if he'd let me. You take a bell band, set it narrow, boost it by ten decibels — deliberately horrible — play the voice, and drag the dot slowly from one-fifty up to five hundred hertz. Somewhere along the way it will suddenly sound awful and honky. That's the problem frequency. You've located it by making it scream.



Corn

Then you flip the boost into a cut.



Herman

Two to four decibel cut, medium width. The over-boost was scaffolding; the cut is the building.



Corn

Ten decibels of ugliness to find four decibels of cleanup. Feels like medicine that has to taste bad.



Herman

Medicine that has to taste loud, briefly. The discipline rule that trips almost everyone, though — and this is the one I'd ask you about, Corn, because you have the producer's instinct of someone who does not want to be here — is that new producers reach for boosts first. They hear a dull recording and they add brightness. Then more brightness. Stack it on top of brightness.



Corn

Which is exactly how a bright room gets a mirror ball.



Herman

Subtractive first. Cut the problem frequencies before you boost anything, because boosting into a compressor amplifies something the compressor is about to squash anyway. There's a structural habit I like for exactly this reason: run ReaEQ twice. First instance does cuts only, before compression. Second instance does the boosts, after compression. You're shaping a stable signal instead of feeding the compressor a moving target.



Corn

Two EQs doing one job each, instead of one EQ negotiating with itself.



Herman

One EQ negotiating with a compressor it can't see. That's worse.



Corn

Okay, compression. Four knobs and a beginner staring at them like they're a cockpit.



Herman

Threshold, ratio, attack, release — plus make-up gain, because compression lowers overall volume, so you bring it back up at the end. Threshold is the line: minus twenty decibels is a common starting point, set so that only the louder words cross it. Ratio is how hard you push back on whatever crosses. Two to one is gentle, four to one is firm. Attack is how fast you clamp — on voice, a few milliseconds to about twenty. On drums the first hit is the transient; on voice it can be a consonant, a sudden laugh, a hard word. Too fast flattens the snap of the consonants, too slow lets the front of loud words through.



Corn

And release is how fast it lets go.



Herman

Roughly eighty to two hundred fifty milliseconds. Too short and you hear it flutter; too long and the quiet words after a loud one stay squashed.



Corn

Here's what I actually want to know, Herman. Every guide gives a range for each knob. But is there one number that matters more than the settings that produce it?



Herman

Yes, and this is the single most useful thing in the whole episode. Gain reduction. Watch the meter and aim for three to six decibels of reduction on the loudest words. Not the knob positions — the result. It reframes the whole thing from "set four knobs correctly" to "watch one number and stop."



Corn

Three to six. That's the escape hatch from knob paralysis.



Herman

And the test is your ear: if you can clearly hear the compressor working — breathing, pumping, the room swelling up in the pauses — back it off. The tell that it's right is that you can't hear it.



Corn

A tool whose success is inaudibility. My kind of equipment.



Herman

One more pro move, for speakers with a wide dynamic range — people who whisper and then laugh at full volume. Two gentler compressors in series, each doing two or three decibels of reduction, almost always sounds more natural than one aggressive compressor doing eight to ten at once.



Corn

Divide the squeeze. I do that with naps and it works.



Herman

I've watched you do it with naps and it works. And I should say the ratio guides disagree — one podcaster guide says four to one is often the most reasonable; another says four to one and above starts sounding processed. Both are defensible. The difference is how much snap you want left in the consonants, which depends on the voice in front of you.



Herman (continued)

These are starting points, not laws.



Corn

So, before we go further — the third tool. The limiter is the one people treat as a second compressor, and that's the whole mistake. And that is the thing about a brickwall limiter, it is protective, it is not a second compressor.

L

Larry

Friends, does your voice suffer from rumble, mud, and honk? Ask your audio professional about EquaCompressil, the only remedy that combines a high-pass filter, a mud cut, and a presence boost in one gentle daily capsule.

L

Larry

EquaCompressil gently sweeps your problem frequencies and turns them into a two to four decibel cut at medium width. Side effects may include sudden honkiness around two hundred hertz, a narrow bell band you cannot stop dragging, unexpected air above ten kilohertz, mild cravings for subtractive equalization, compression that reacts to your rumble, a lookahead delay of a few hundred microseconds, and in rare cases a fast-food drive-through speaker where your radio voice used to be. Do not stop taking EquaCompressil without consulting your limiter. EquaCompressil. Because knowing when to stop is the real skill.



Herman

I already have some of that, actually. The second bottle. The presence boost was fine but the cravings for subtractive EQ are real — I caught myself cutting frequencies on a voice memo Hannah's cousin sent, just for the calm of it.



Corn

Anyway. The limiter.



Herman

Yes. Sound On Sound's technical editor, Hugh Robjohns, puts the role plainly: a limiter is fundamentally protective — it exists to make sure whatever its output is fed into never receives excessive signal levels. On a voice track it should be almost invisible. If it's working hard, the answer is upstream: back off the compressor, lower the tone boosts. A limiter earning its keep is one you never hear.



Corn

Mechanically, how does it catch a peak that hasn't happened yet?



Herman

Lookahead. It analyzes the signal before outputting it, which means inserting a tiny delay — a few hundred microseconds — so the attenuation lands just before the transient arrives instead of just after. And the compressor-limiter distinction is a matter of degree, not kind. A limiter is a compressor with a very high ratio: ten to one up to infinity to one, and most real ones sit around twenty to one. At that ratio, anything over the threshold comes out at the threshold, near enough. A wall.



Corn

Where does the ceiling go?



Herman

Between minus one and minus a tenth of a decibel below full scale, to leave margin for MP3 and AAC encoding — and true-peak mode guards against clipping at digital-to-analog conversion, which is where intersample peaks sneak in and bite you.



Corn

Now the numbers people argue about in forums at two in the morning. Loudness.



Herman

The confusion is worth a whole beat, because it's muddled. Apple Podcasts recommends about minus sixteen LUFS integrated, plus or minus one, with true peak at or below minus one dBTP — and one figure for mono and stereo alike. The minus nineteen LUFS mono number that floats around is an industry convention, not Apple's spec. And here's the physics of why people get confused: a dual-mono file measures about three LU louder than expected, because the loudness meter sums correlated channels. Two identical channels read as more energy.



Corn

So someone measures their mono file, sees it three units hot, concludes the target must be different, and a folklore is born.



Herman

That's exactly how it happens. Spotify's minus fourteen figure comes from its documentation for music tracks, not a podcast requirement. There's no single enforced podcast standard — Apple's is the one worth aiming at because it's the biggest doorway.



Corn

The doorway measures your file on the way in.



Herman

Gently. The platform will turn it down, not reject it.



Corn

REAPER itself. What does the stock arsenal actually cover?



Herman

Everything in the chain, which is the quiet scandal of it. ReaLimit came in during the six-point-x cycle as a brickwall limiter and loudness maximizer, multiple lookahead modes, visual peaks display. Version seven added per-track LUFS metering and render



Herman (continued)

statistics with a shareable HTML loudness and clipping report — so you can check your minus sixteen without buying anything. And the full voice chain, five stock plugins, six slots: ReaFir in subtract mode for noise — you build a profile from a few seconds of room tone, then untick the auto-build — then ReaEQ for the high-pass and cuts, then ReaXcomp for sibilance, then ReaComp for the broad level control, then ReaEQ again for the boosts, then ReaLimit last. Costs nothing beyond the license you already have.



Corn

You said ReaXcomp handles sibilance, but REAPER has no dedicated de-esser, doesn't it? That's the gap that sends people shopping.



Herman

It does send people shopping, and it needn't. ReaXcomp is a multiband compressor — set up one band covering about five kilohertz to nine, fast attack, low threshold, and it clamps sibilance only when sibilance happens. That's the entire argument for a de-esser over an EQ cut: a cut at that frequency would dull every word all the time; the dynamic band only turns down when an S actually occurs. A second band at two hundred to four hundred with gentler settings handles boxiness the same way.



Corn

Two problems, one plugin, no purchases.



Herman

And ReaComp has two features nobody uses that fix real problems. The detector input has its own filtering — you can tell the compressor to ignore low frequencies when deciding whether to squeeze. That stops every plosive from grabbing the whole voice. And auto make-up gain keeps the output level steady while you move the threshold around, which means while you experiment you're comparing tone rather than loudness.



Corn

Louder always sounds better. That's the oldest trick in the ears.



Herman

Louder sounds better, so if the compressor makes you quieter while you tune it, you'll undo your own good decisions to get the volume back.



Corn

Herman, I want to ask about the addictive part, because Daniel said it and I don't think he was joking. He said he never would have guessed that dragging points along a curve could be weirdly addictive. Do you buy that it's the interaction itself, not the sound improvement?



Herman

Completely. The mechanics are the same everywhere — left and right moves the frequency, up and down sets the gain, the mouse wheel over a dot changes the width, double-click flattens the band. Roon, announcing their MUSE EQ, warns in their own copy that the controls are highly interactive and — their word — highly addictive. A company warning you off its own wheel.



Corn

That's a cigarette packet with a compliments slip inside.



Herman

And Mastering The Mix released a whole EQ this year built around it — you pull a frequency towards you to make it louder, push it away to make it quieter, the room expands as you work. Their pitch is that it feels less like a tool and more like an instrument.



Hilbert

It adjusts the width. Whoever decided the wheel over the dot should do gain instead of width was not thinking about anyone's hand but his own. Herman, you've been saying "mouse wheel changes Q" like it's settled law — it isn't, it's a convention, and it's the wrong one.



Herman

You'd rather the wheel did width and gain stayed on the drag?



Hilbert

I'd rather the wheel did what the wheel does on a fader, which is the fine version of the thing you're already doing. Corn dragged the wheel on that de-esser band earlier like a man stirring soup — sorry, Corn, that cough is coming out of your last thirty seconds, I'll cut it in the edit — and the gain jumped half a decibel at a time. Width wants to be nudged. Gain wants to be dragged. That's the whole design and somebody inverted it.



Corn

Half-decibel increments and I still found the sibilance.



Hilbert

You did, eventually. My cousin sat with me one evening while I had a ReaEQ open on a voice track — hers, a recording she'd made of a eulogy, nothing to fix on it, the room was good. I spent the whole evening dragging one point back and forth from two hundred to about five hundred hertz without the audio playing once. Just watching the curve move. She left around eleven. I kept going a while after. I've still got a screenshot of the curve from that night. A cut at three hundred and ten, medium bell, and the curve had a shape to it I've never got again. It's been my desktop wallpaper since. Fourteen months.



Herman

You never listened to the result of the cut?



Hilbert

There wasn't a result. I flattened the band before I closed the session. The curve was the point.



Herman

The image of the shape, not the effect of it.



Hilbert

The effect was fine too, I'm sure. Corn's point about it being stirring soup — no. He's right that it's slow. But slow on the wheel is how you land a narrow bell where you mean it, and I'd take slow and accurate over fast and half a decibel off. My cousin still asks whether I fixed her recording. I tell her it didn't need fixing.



Corn

So the curve is worth watching even when there's nothing wrong with the audio.



Herman

Which is the exact tension in the working-engineer world, and it's the best argument in the whole subject. On Gearspace, a guy with a couple of decades of in-the-box work described the opposite epiphany — he now prefers knob-only parametric EQs for fine-tuning, over graphic EQs where you see the curves, because the visual feedback seduced him. Watching the FFT move in real time, he stopped listening as critically, started using his eyes more.



Corn

He traded the picture for the sound because the picture was winning.



Herman

And others in the same thread pushed back — the curve shows boosts and cuts accurately, it speeds up the workflow, it aids learning. Both are true. The curve is a great teacher and a bad crutch, and it's the same curve.



Corn

I'd land it like this: the visual is how you learn where three hundred hertz lives. It's not how you learn whether your voice needs a cut there.



Herman

Yes — the curve answers where, never whether. Whether is ears. And that's a habit, not a setting: look at the curve to aim, then close your eyes and drag, and see if you still believe the move when you hear it.



Corn

The other thing the two-EQ habit does, cuts before boosts, compressor in the middle — it's not a recipe, it's a reasoning. Every piece of it is "what does the next thing in line react to."



Herman

That's the refresher I'd give Daniel. Not a preset, a set of questions. What's the compressor reacting to? What's the limiter catching? Where does the gain reduction meter sit on the loud words — three to six, or breathing and pumping? And the loudest-number rule of the whole episode: if you can hear any of it working, back it off.



Corn

Used lightly, radio. Used heavily, drive-through. And he found the fun part on his own — the dot, the drag, the curve answering back.



Herman

Herman Poppleberry, thanking our producer, Hilbert Flumingtop, at the desk as always. If today's refresher earned its ten decibels of honk, leave us a review wherever you get the show — it helps.



Corn

If you enjoyed this one, try episode seven, Building Custom ASR Tools. This has been My Weird Prompts. We'll be back soon.

SOURCES

What the research for this episode read before the script was written. Primary sources first.

1. **REAPER official site (Cockos), feature list and REAPER 7 changelog summary, retrieved 2026-11-16. [primary]**
   https://www.reaper.fm/

2. **Acast, A guide to vocal compression for podcasters, retrieved 2026-11-16 (no publish date shown).**
   https://www.acast.com/en/blog/a-guide-to-vocal-compression-for-podcasters

3. **Jeff Pierson, Journalism Training, EQ and Compression on a Voice (Lesson 4 of 7), retrieved 2026-11-16.**
   https://jeffapierson.com/journal_tools/training/08-audio-production/04-eq-and-compression.html

4. **Jeff Pierson, Journalism Training, Loudness and Delivery (Lesson 7 of 7), retrieved 2026-11-16 (via search snippet).**
   https://jeffapierson.com/journal_tools/training/08-audio-production/07-loudness-and-delivery.html

5. **Pristine Podcasts, Podcast EQ and Compression Settings That Actually Work, published 2025-05-22, retrieved 2026-11-16.**
   https://pristinepodcasts.com/blog/podcast-eq-compression-settings

6. **Cedar Sound Studios, Reaper Vocal Chain Setup: 5 Essential Plugins for Clean Vocals, retrieved 2026-11-16.**
   https://www.cedarsoundstudios.com/blogs/news/reaper-vocal-chain-setup-5-essential-plugins-for-clean-vocals

7. **The REAPER Blog, Dialog FX Chain with stock plugins, published 2020-06-05, retrieved 2026-11-16.**
   https://reaper.blog/2020/06/dialog-fx-chain/

8. **Sound On Sound, Q. Is there more to a Limiter than a Compressor with an infinite ratio? (Hugh Robjohns), retrieved 2026-11-16.**
   https://www.soundonsound.com/sound-advice/q-there-more-limiter-compressor-infinite-ratio

9. **iZotope, An introduction to limiters (and how to use them), retrieved 2026-11-16.**
   https://www.izotope.com/community/blog/an-introduction-to-limiters-and-how-to-use-them

10. **Peak Studios, Limiter**
    compressor in the mix, mastering, retrieved 2026-11-16.
    https://www.peak-studios.de/en/limiter/

11. **Critical Listening Lab, Podcast Loudness Standards**
    LUFS for Apple Podcasts & Spotify, retrieved 2026-11-16.
    https://www.criticallisteninglab.com/en/learn/loudness/podcast

12. **Hanna Eng, Podcast Loudness: What LUFS to Target, updated 2026-10-16, retrieved 2026-11-16.**
https://www.hanna-eng.com/guides/podcast-loudness/

13. **SoundForgePro, Podcast Loudness Standard: LUFS & True Peak, updated 2026-10-22, retrieved 2026-11-16 (via search snippet).**
https://soundforgepro.com/podcast-loudness-standard/

14. **Avid, How to EQ Vocals: Essential Tips for Clarity & Presence, retrieved 2026-11-16 (via search snippet).**
https://www.avid.com/pro-tools/learn/vocal-mixing/vocal-eq

15. **Orphiq, Vocal EQ Cheat Sheet, published 2026-07-23, retrieved 2026-11-16 (via search snippet).**
https://orphiq.com/resources/vocal-eq-cheat-sheet

16. **Roon Labs, Introducing MUSE: Precision Audio Control on the go, retrieved 2026-11-16 (via search snippet).**
https://blog.roonlabs.net/introducing-muse-precision-audio-control-on-the-go/

17. **Musiques Live, Mastering The Mix releases Mixvault, published 2026-08-04, retrieved 2026-11-16 (via search snippet).**
https://www.musiques.live/2026/08/04/mastering-the-mix-releases-mixvault-vst3-au-aax-2/

18. **Gearspace forum thread, Epiphany: I think I prefer parametric EQs!, retrieved 2026-11-16 (snippet only; Cloudflare-blocked).**
https://gearspace.com/threads/epiphany-i-think-i-prefer-parametric-eqs.1347842/

19. **Audio Compressor, Audio Compressor Attack and Release Explained, retrieved 2026-11-16 (via search snippet).**
https://audio-compressor.org/audio-compressor-attack-release/

20. **Kilohearts, Carve EQ documentation, retrieved 2026-11-16 (via search snippet).**
https://kilohearts.com/docs/carve_eq