

VCA Comp

VCA COMPRESSOR · AU + VST3 · macOS

ZEROPINK

USER MANUAL



Every control the circuit can give you, and nothing decided on your behalf. The precise one.

We all want the perfect mix.

Before anything else: get to ZERO

01

VCA Comp assumes you have set your level before you start compressing. That is not housekeeping — it is the whole method, and every ZEROPINK plugin is built around it.

0 VU = -18 dBFS. That is the reference the meters are calibrated to, and it is where analogue equipment expects to be fed. Modern digital sessions routinely run 15–20 dB hotter than this, which is why compressors that sound effortless on one engineer's mix sound crushed and lifeless on another's. Same plugin, different starting level.

On a VCA compressor the consequence is specific: **THRESHOLD is an absolute level**, not a relative one. A threshold of -18 dB means -18 dB, and if you arrive 12 dB hot, that setting is suddenly 12 dB deeper into the signal than you intended. Every stored preset, every threshold you have learned by feel, is only meaningful from a known starting point.

1. Play the loudest section of the part.
2. Bring the INPUT needle to around 0 VU. Two ways to get there:

Clip gain / region gain in your DAW — the better fix. It corrects the level at source, so every plugin downstream inherits a sane starting point.

INPUT GAIN on the plugin — the quick fix, and the right one when the level is correct but the part simply needs trimming into the detector.

3. Only then set THRESHOLD.

The STAGE IN readout underneath tells you exactly how far off you are and what to do about it. It does not report a number and leave you to interpret it — it names the state and prescribes the correction:

STAGE IN **-5.0** VU LOW **push input +5.0**

That reads: you are 5 dB below reference, and the fix is to add 5 dB of input gain. When you land, the state word changes and the prescription disappears.

Get this right and the rest of the plugin behaves predictably from session to session. Skip it and you will spend the afternoon fighting a threshold that keeps meaning something different.



The guidance strip. Three states, three prescriptions — read left to right as you work.

What VCA Comp is

02

A digital model of a bus-compressor-class VCA design. Gain reduction is performed by a voltage-controlled amplifier: a control voltage derived from the sidechain sets the gain, directly and linearly. Nothing in the path has to warm up, saturate or relax.

That linearity is the character — or rather, the deliberate absence of one. Where the optical cell has a memory and the FET has a temper, the VCA does exactly what the control voltage tells it to, which means **every timing decision is yours**. Attack, release, threshold, ratio and the sidechain filter are all exposed, and the compressor has no opinion of its own to override them.

This is why VCA designs became the standard mix-bus compressor. Predictable, repeatable, and quiet enough to leave across a whole record without accumulating character you did not ask for.

It is also the most demanding of the three. Six controls that interact means six ways to be wrong, and the plugin will not save you from a bad setting the way a slower topology might.

GOOD AT

Mix bus and stem glue. Drum busses that need to breathe with the tempo. Anything where you want a specific attack and release rather than a flavour. Programme material where transparency matters and the compression should be felt as cohesion, not as an effect.

NOT GOOD AT

Colour. If you want the part to sound compressed — thick, forward, obviously worked — this is the wrong tool and no setting will get you there. That is FET territory. Equally, if you want a compressor that smooths a performance without you deciding how, OPTO does that better by having fewer choices.

Choose VCA Comp when you know what you want the envelope to do. Choose one of the others when you want the circuit to decide.

The three meters

03

All three are true VU — around 300 ms integration, IEC/ASA ballistics. They show what you would feel, not sample peaks. On a compressor this controllable that matters: the needle and your attack setting are measuring different things, and reading one as the other is the commonest way to mis-set this plugin.



INPUT — level arriving at the detector and the VCA, after INPUT GAIN. Set this to 0 VU first; the threshold you choose next only means what you think it means from here.



GAIN REDUCTION — how much gain the VCA is removing, reading right to left from 0. On a well-set VCA this needle should move like the music, not flicker.



OUTPUT — level leaving the plugin, after OUTPUT GAIN. Bring this back alongside the input needle so you judge the compression, not the level.

INPUT and OUTPUT share one scale — -20 · -10 · -7 · -5 · -3 · 0 · +3 — with the stretch above 0 marked in pink. That region is not a warning light; it is the working headroom above reference, and material genuinely lives there. Sustained residence there means your gain staging has drifted. GAIN REDUCTION carries its own scale — 20 · 15 · 10 · 5 · 0 — reading right to left, so the needle rests at 0 and travels left as reduction deepens.

INPUT GAIN · THRESHOLD

INPUT GAIN · ±24 dB

Trims the level arriving at the detector and the VCA. Unlike FET Comp, this is *not* the compression control — it is the staging control. Use it to reach 0 VU and then leave it alone.

The knob has a detent at unity: it snaps to exactly 0.0 as you pass it, and the engraved zero at the centre lights pink when it lands. Double-click returns it to unity from anywhere.

THRESHOLD · -48 ... 0 dB

The level at which compression begins. Above it the ratio applies; below it the plugin is doing nothing. This is the control that decides *how much* of the material gets compressed, as distinct from ratio, which decides how hard.

Set it by watching the GAIN REDUCTION needle rather than by number. Lower it until the needle moves the amount you want on the loudest passages, then check a quiet passage to make sure it is not working there too.



Reading the prescription. STAGE IN names your position against reference in three states:

STATE	MEANS	PRESCRIPTION
LOW	Below reference — your threshold will read deeper than it is	push input +n
ON TARGET	At reference. Threshold settings now mean what they say	✓ staged
HOT	Above reference — every threshold is deeper than you intend	trim input -n

The number beside the prescription is the exact correction, so there is nothing to calculate.

A useful discipline: once STAGE IN says staged, do not touch INPUT GAIN again for that part. If you want more compression, lower the threshold. Keeping the staging fixed is what makes your threshold settings portable between sessions — and what lets a preset mean anything at all.

ATTACK · 0.1 – 100 ms

How fast the VCA responds once the signal crosses threshold. This range is enormous compared to a FET — three orders of magnitude — and that breadth is the point. You are not picking a flavour, you are deciding precisely which part of the envelope gets compressed.

Fast (0.1–3 ms) catches transients and removes them. Medium (5–20 ms) lets the initial hit through and compresses the body behind it. Slow (30–100 ms) leaves the attack entirely intact and works only on sustain — on a mix bus this is usually what you want, because it preserves the punch that makes the record feel alive.

RELEASE · 50 – 2000 ms, or AUTO

How quickly gain returns once the signal falls back. Short release re-opens between events and pulls up everything in the gaps — density, room, pump. Long release stays down through a phrase and sits transparently behind the source.

On a bus, release decides whether the compression breathes with the tempo or fights it: set it so the needle returns to rest just before the next downbeat and the compressor locks to the groove.

AUTO RELEASE

Drag RELEASE past the top of its travel and it latches to **AUTO** — the ring fills completely and the readout in the knob reads the word instead of a figure.

In AUTO the recovery time is derived from the programme itself: it moves quickly after a short transient and slowly after a sustained passage, continuously, without you choosing a compromise between the two. It is the one place on this plugin where the circuit makes a decision for you.

Reach for AUTO on dynamic programme material — a full mix, a live take, anything where a fixed release is either too fast for the quiet parts or too slow for the loud ones. Use a fixed value when you want the compression to have a rhythm of its own. Dragging back down leaves AUTO; double-click resets to 200 ms.

Starting points. Not rules — places to begin listening from:

SOURCE	ATTACK	RELEASE	WHY
Mix bus	30	AUTO	Keeps transients; lets the programme set recovery
Drum bus	10 – 30	100 – 250	Time the release to the tempo and it pumps in time
Vocal	3 – 10	150 – 400	Catches the line without chewing consonants
Bass	10 – 20	400 – 800	Long release avoids modulating the fundamental
Strings	30 – 100	AUTO	Slow attack keeps bow noise; AUTO tracks the phrasing

Five caps: 1.5 · 2 · 4 · 8 · 10. Ratio decides how much of every dB past the threshold is held back — at 4:1, four dB in becomes one dB out. Where threshold decides *how much* of the signal is compressed, ratio decides *how hard*.

RATIO	CHARACTER	REACH FOR IT ON
1.5	Barely a compressor — a gentle tilt in dynamics	Mastering, mix bus, wide-threshold glue
2	The classic bus setting — cohesion without grip	Mix bus, stems, orchestral submixes
4	Even-handed and audible; the general-purpose choice	Vocals, bass, drum busses, most single tracks
8	Firm control — holds a part in place	Uneven performances, aggressive vocal, snare
10	Effectively limiting — a ceiling rather than a shape	Peak control, holding a part absolutely still

Ratio and threshold trade against each other. A high ratio with a high threshold compresses only the peaks, and hard — the same average reduction reached with a low ratio and a low threshold compresses nearly everything, gently. Those two settings measure the same on the GR needle and sound nothing alike.

On a bus, prefer the second approach: 1.5 or 2, threshold low enough that the needle is always doing a little. Reserve high ratios for individual tracks that genuinely need holding down.

HP FILTER · 20 – 500 Hz

A high-pass filter on the **detector only** — not on the audio. Nothing is removed from the signal you hear; the filter changes what the compressor *listens* to.

This exists because low frequencies carry most of the energy in a mix. A kick drum and a bass note contain far more power than a vocal, so a full-band detector reduces gain on the whole mix every time the kick lands. The result is a mix that ducks and pumps in the midrange for reasons that have nothing to do with the midrange.

Raise the filter until the GR needle stops following the kick. The compressor keeps working on everything above it, and the low end passes uncompressed.

SETTING	EFFECT
20 Hz	Effectively off — the detector hears everything
60 – 90 Hz	Removes sub weight from the sidechain; keeps the kick's body
100 – 150 Hz	The usual mix-bus setting — kick and bass stop triggering
200 – 500 Hz	Detector responds mostly to midrange and vocals

This is the single most useful control on a mix bus compressor and the one most often left at 20 Hz. If bus compression is making your mix feel smaller rather than more cohesive, come here first.

Reading REDUCTION - OUTPUT GAIN

The REDUCTION readout names how hard the compressor is working and tells you which way to go:

REDUCTION **-4.5** dB **RIDE** classic 3-6 ✓

STATE	REDUCTION	WHAT IT MEANS	TYPICAL USE
KISS	0 to -3 dB	Barely engaging — peaks only	Mastering, gentle bus glue, safety
RIDE	-3 to -7 dB	The classic working range — the panel calls it <i>classic 3-6</i>	Mix bus, drum bus, vocals, most work
DEEP	-7 to -12 dB	Audibly compressing — you hear it working	Uneven performances, deliberate pump
CRUSH	-12 dB +	Heavy reduction, the effect is the point	Parallel blends, effect sends

The prescription beside the state tells you what to do to move between them. On this plugin there are two routes down — raise the threshold, or drop a ratio — and they sound different. Raise the threshold when the compression is reaching too far into the material; drop the ratio when it is reaching the right distance but pressing too hard.

On a bus, CRUSH is almost always a mistake rather than a choice. If you want that much reduction on a mix, use MIX to blend it back rather than committing to it.

OUTPUT GAIN

Makeup gain after the VCA. Because everything in ZEROPINK is referenced to 0 VU, the goal is not “as loud as possible” — it is to return to reference so you can judge the compression rather than the level increase.

STAGE OUT prescribes in both directions and confirms when you arrive:

LOW

trim output +4.5

-4.5 VU

HOT

trim output -3.1

+3.1 VU


ON TARGET

IN/OUT matched ✓

-0.2 VU

IN/OUT matched means input and output are at the same level, so bypassing the plugin is an honest comparison. Louder always sounds better for the first few seconds; this removes that illusion from the decision. On a transparent compressor it is the only reliable way to tell whether it is helping at all.





OS LIVE / OS RENDER — oversampling for the nonlinear stage, set independently for playback and for export.

	OPTIONS	DEFAULT
OS LIVE	Off · 2x · 4x · 8x · 16x	4x
OS RENDER	Off · 2x · 4x · 8x · 16x · 32x · 64x · 128x	16x

Oversampling runs the nonlinear stage at a multiple of the session rate so the harmonics it generates land where they belong instead of folding back down into the audible range as aliasing. More is cleaner, and more costs CPU.

A VCA generates less harmonic content than a FET stage, so it needs less help — but it is also the compressor most likely to sit on the mix bus, where the whole record passes through it. What little aliasing it produces lands on everything.

The split exists so you never have to choose between a responsive session and a clean master. Work at 4x — low latency, light CPU, and audibly fine — and let the render setting take over on export, where the extra passes cost you nothing but time. The render figure is the one that reaches the listener.

Off is there for CPU emergencies and for null-testing. The high render settings are for the stubborn cases: dense material with a lot of high-frequency content driven hard, or a fast attack working on bright programme. Returns diminish steeply past 16x, so treat 32x and above as a tool for a specific problem rather than a default worth paying for.

GLOBAL — applies the oversampling setting across all ZEROPINK plugins in the session, not just this instance. Switch one to 16x for a render and every ZEROPINK plugin in the project follows, so you are not hunting through inserts before a bounce.

Why oversample at all

09

THE PROBLEM: ALIASING

A compressor is a *nonlinear* process. It does not merely make the signal quieter — it multiplies the signal by a gain that is itself changing, and any nonlinearity creates frequency content that was not in the input. New harmonics appear above the original material.

A digital system can only represent frequencies below half its sample rate — the Nyquist limit. At 44.1 kHz that ceiling is 22.05 kHz. Harmonics generated above it cannot simply exist and be inaudible, as they would in an analogue circuit. They **fold back** down into the audible band, mirrored around Nyquist.

A 15 kHz tone driven hard produces a 3rd harmonic at 45 kHz. At 44.1 kHz that folds back to about **0.9 kHz** — it reappears near 900 Hz, in the middle of the vocal range, at no musical interval to anything playing.

That is the whole problem. Harmonic distortion sits *above* the note and sounds like warmth. Aliased distortion sits at arbitrary frequencies *below* it and sounds like grit, glare or a metallic edge, and it moves the wrong way when the music moves, so the ear never files it as part of the instrument.

INTERSAMPLE BEHAVIOUR

Samples are not the waveform. They are points on it, and the reconstructed analogue signal passes *between* them — often overshooting the highest sample by a decibel or more. These intersample peaks are invisible to a detector that only reads sample values.

This matters twice over:

Detection. A sidechain running at base rate can miss a peak that occurs between two samples, so the compressor responds late or not at all. With attack times down to 0.1 ms available here, you can set a response faster than the sample grid can describe — and then not get it.

Modulation. Applying gain is multiplication, and multiplying two signals produces sum and difference frequencies — sidebands around every partial. A gain signal that changes abruptly on a sample grid throws those sidebands high, where they alias. A fast attack on a bright source is exactly that condition.

Running the detector and the gain stage at a higher rate lets the gain curve resolve smoothly in time, so its sidebands stay in band and the transient is actually seen. On a VCA — where you may have asked for a genuinely fast attack — this is what makes the setting mean what it says.

WHAT OVERSAMPLING DOES

The plugin raises the sample rate before the nonlinear stage, processes there, filters, and returns to the session rate:

1. **Upsample** — insert samples and interpolate.
At 4x the Nyquist ceiling moves from 22 kHz to 88 kHz.
2. **Process** — detection and gain now run with room above them, so new harmonics land where they belong instead of folding.
3. **Filter** — a steep low-pass removes everything above the original Nyquist.
4. **Downsample** — return to session rate with nothing left to fold.

AUDIO INTEGRITY

Aliasing is not reversible. Folded content lands *inside* the audible band, on top of the music, at no fixed relationship to it — no later EQ, filter or restoration can separate it out. Unlike noise or a bad EQ move, it cannot be undone in the next session.

That is why the render setting is the one that matters, and why it is worth paying for on export even when you never hear the difference while working.


Staging helps too. Alias energy grows with how hard the nonlinearity is driven. Arriving at 0 VU means the compressor is working in the range it was designed for rather than being pushed there by a hot session.

The fold, drawn

10

A single 15 kHz tone, driven hard enough to generate harmonics. What happens to them depends entirely on how much room there is above the music.

A · AT THE SESSION RATE — 44.1 kHz, NO OVERSAMPLING



Neither harmonic can exist above Nyquist, so each is **mirrored back** around it — 30 kHz reappears at 14.1 kHz, 45 kHz at 0.9 kHz. Note where they land: the 3rd harmonic, which belongs an octave and a fifth *above* the source, now sits four octaves *below* it. Nothing about that interval is musical, and if the player bends the note upward the folded partials slide *down*.

B · THE SAME SIGNAL AT 4× — 176.4 kHz INTERNAL, NYQUIST AT 88.2 kHz



With Nyquist four times higher the harmonics have room to exist as harmonics. The anti-alias filter removes them before the rate drops, so the signal returns to 44.1 kHz carrying only what belongs there.

Why it must be avoided: the pink bars in plot A are inside the audible band, on top of the music, at no fixed relationship to it. No filter, EQ or restoration can separate them out afterwards — they are indistinguishable from programme material. The only place to prevent them is here, before the render.

MIX — parallel blend. 100% is fully compressed. Lower values blend the uncompressed signal back in, giving you New York-style parallel compression without a second bus.


On a transparent compressor this control has a particular use: it lets you commit to a heavy setting and then dose it. A 10 dB reduction at 40% mix holds the level far more gently than a 4 dB reduction at 100%, and keeps the transients the fast setting removed.

0 VU REF — the reference calibration. Five settings: -18 · -21 · -24 · -27 · -30 dBFS.

-18 is the default and the right choice for programme material and mix-bus work. The lower settings are for percussive sources with a high crest factor: a snare or kick has peaks far above the average the VU reads, so at -18 the needle can sit politely near 0 while the actual peaks are already close to the ceiling. Dropping to -24 or -27 pushes the whole scale down and gives those transients room above 0 VU where they belong.

Changing this moves every meter's zero point. It does not change the audio — it changes what the meters call zero, and therefore what the STAGE readouts prescribe. It does not move the THRESHOLD, which remains an absolute level.

SKIN — switches the faceplate between Light and Dark mode. Cosmetic only; it has no effect on audio or on metering behaviour. Pick whichever holds up in your room — dark rooms and long sessions usually favour Dark, bright rooms Light.



One design, two skins. Only the metal changes — glass, meters and the pink ramp carry over.

Working method

12

A sequence that uses the plugin the way it was designed. The order matters more here than on the other two, because every control affects what the next one does:

1

Stage in. INPUT GAIN until the INPUT needle sits at 0 VU and STAGE IN stops prescribing. Then leave it alone — this is what makes your threshold mean something.

2

Set the sidechain filter. On anything with low-end energy, raise HP FILTER until the GR needle stops following the kick. Do this before the timing, not after.

3

Pick a ratio. 1.5 or 2 for bus work, 4 for single tracks, 8 or 10 to hold something still. Choose the behaviour first.

4


Lower the threshold until REDUCTION reads RIDE — the panel's *classic* 3–6. Check a quiet passage as well as a loud one.

5

Set the timing. ATTACK decides whether the transient survives; RELEASE decides whether the compression breathes with the tempo. Try AUTO on dynamic material.

6

Match and compare. OUTPUT GAIN until STAGE OUT reads IN/OUT matched, then bypass. If you cannot hear an improvement, you do not need the compressor on this source.



The trial countdown sits in the header and disappears the moment a licence is entered.

VCA Comp installs as a fully functional 30-day trial. Nothing is disabled, nothing is muted, no noise is injected. A header banner shows the time remaining.

When the trial ends the plugin outputs silence and displays a persistent banner. Your session, settings and automation are preserved — entering a licence restores audio immediately, exactly where you left off.

Entering your licence. Purchase delivers a signed licence block by email. Paste it into the plugin's licence dialogue. Verification happens offline on your machine; the plugin never contacts a server, and it will keep working forever on a machine with no internet connection.

Deauthorize removes the licence from the current machine, returning it to trial state, so you can move it to another. There is no activation limit and no account to log into.

Lost your licence? Contact support with your order details and it will be reissued.

TROUBLESHOOTING

The plugin does not appear in my DAW. Restart the DAW after installing. Logic rescans on launch and will validate the AU the first time it sees it. If it still does not appear, check that the DAW is not running in Rosetta.

My mix pumps in the midrange. The detector is following the kick. Raise HP FILTER to 100–150 Hz.

The threshold feels wrong on every preset. Check STAGE IN. Threshold is an absolute level, so a hot input makes every stored threshold deeper than intended.

RELEASE shows AUTO and will not move. It has latched to the AUTO detent past the top of its travel. Drag downward to leave it, or double-click to reset to 200 ms.

The mix sounds smaller with the compressor in. Match the output first — an unmatched comparison is meaningless. If it still sounds smaller, lower the ratio and the threshold together.

The trial banner says expired but I purchased. The licence has to be entered in the plugin — purchasing does not unlock it automatically, because verification is offline.

Specifications

Formats	AU · VST3
Platform	macOS, Apple Silicon
Topology	VCA, variable threshold. Ratio 1.5 · 2 · 4 · 8 · 10
Threshold	-48 ... 0 dB
Attack	0.1 – 100 ms
Release	50 – 2000 ms, or AUTO (programme-dependent)
Sidechain HP	20 – 500 Hz, detector only
Gain range	Input and output ±24 dB, detent at unity
Metering	True VU, ~300 ms integration, IEC/ASA ballistics
0 VU reference	-18 dBFS default · -21 / -24 / -27 / -30 dBFS selectable
Oversampling	Live: Off-16× (4× default) · Render: Off-128× (16× default)
Sample rates	Validated 11.025 – 192 kHz

LATENCY

Reported to the host at every setting, so delay compensation is exact. The figure in samples is fixed; the time it represents shortens as the session rate rises.

OS LIVE	SAMPLES	44.1 kHz	48 kHz	96 kHz	192 kHz
Off	0	0.000 ms	0.000 ms	0.000 ms	0.000 ms
2×	49	1.111 ms	1.021 ms	0.510 ms	0.255 ms
4× DEFAULT	61	1.383 ms	1.271 ms	0.635 ms	0.318 ms
8×	65	1.474 ms	1.354 ms	0.677 ms	0.339 ms
16×	67	1.519 ms	1.396 ms	0.698 ms	0.349 ms

At the 4× default this is well under 1.5 ms at any rate — low enough to track through. Set OS Live to Off if you need a genuinely zero-latency path while recording.