

FET Comp

FET COMPRESSOR · AU + VST3 · macOS

ZEROPINK

USER MANUAL



Fast, forward and deliberately unsubtle. Drive the input, pick a ratio, get out of the way.

We all want the perfect mix.

Before anything else: get to ZERO

01

FET 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 FET compressor this matters more than on any other topology, because **INPUT GAIN is the compression control**. There is no threshold knob to compensate with. If you arrive 12 dB hot, you are already deep into reduction before you have touched anything, and every ratio will sound like the wrong one.

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 — to reach reference, and then to drive the FET stage past it once the staging is right.

3. Only then pick a ratio.

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 **-1.7** VU LOW **push input +1.7**

That reads: you are 1.7 dB below reference, and the fix is to add 1.7 dB of input gain. When you land, the state word changes and the prescription disappears.

Reaching 0 VU is the starting line, not the finish. On this plugin you then push *past* it deliberately — but from a known place, so you can hear what the drive is doing rather than discovering it.



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

What FET Comp is

02

A digital model of a 1176-class FET compressor. In the hardware it emulates, gain reduction is done by a field-effect transistor acting as a voltage-controlled resistor. A transistor has no mass and no thermal lag, so it can clamp in microseconds — orders of magnitude faster than a photocell.

That speed is the entire character. FET Comp catches the front edge of a transient rather than arriving after it, and the act of clamping that fast generates harmonic content of its own. The distortion is not a side effect to be minimised; it is why engineers reach for these units.

It has no threshold control. The threshold is fixed in the circuit. You do not lower a threshold to the signal — you raise the signal to the threshold with INPUT GAIN, and make the level back up afterwards with OUTPUT GAIN. Those two knobs, worked against each other, are the compressor.

Ratio is a button bank, not a continuous knob, because on the original it was four physical switches. Engage all of them at once and the circuit does something it was never designed to do — see the next page.

GOOD AT

Drums — snare, kick, room mics, whole kits. Rock and pop vocals that need to sit forward and stay there. Bass that has to be felt at a constant level. Electric guitar. Anything where you want to *hear* the compression and want the attack shortened.

NOT GOOD AT


Transparency. If the brief is “hold the level without anyone noticing”, this is the wrong tool — that is OPTO territory. FET Comp on a delicate acoustic performance or a finished mix will usually be too fast, too obvious and too coloured, and no amount of gentle setting hides it.

A useful test: if you find yourself trying to make FET Comp disappear, you want a different compressor. Use it when the colour is the point.

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 fast that gap is worth remembering: the FET is acting on events the needle is too slow to draw, so you read the meters for level and judge the speed by ear.



INPUT — level arriving at the FET stage, after INPUT GAIN. This is the one you set first, and then the one you push deliberately.



GAIN REDUCTION — how much gain is being removed, reading right to left from 0. Expect it to snap rather than glide; that is the topology behaving correctly.



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

INPUT and OUTPUT share one scale — $-20 \cdot -10 \cdot -7 \cdot -5 \cdot -3 \cdot 0 \cdot +3$ — with the stretch above 0 marked in pink. That region is not a warning light; it is the working headroom above reference, and on this plugin you will visit it on purpose. GAIN REDUCTION carries its own scale — $20 \cdot 15 \cdot 10 \cdot 5 \cdot 0$ — reading right to left, so the needle rests at 0 and travels left as reduction deepens.

The controls · INPUT GAIN

This is the compression control. Because the threshold is fixed, the only way to ask for more reduction is to send more signal. INPUT GAIN therefore sets depth, and along with it the amount of harmonic colour the stage contributes.

Use it in two moves. First reach 0 VU so STAGE IN stops prescribing. Then push beyond it and watch the GAIN REDUCTION needle until the amount of reduction is what you want to hear.

Every dB you add here is a dB you will take back with OUTPUT GAIN. That is the trade the plugin is built around, and it is why the two knobs are the same size at opposite ends of the panel.



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

STATE	MEANS	PRESCRIPTION
LOW	Below reference — the FET stage is being underfed	push input +n
ON TARGET	At reference. Start here, then drive past it deliberately	✓ staged
HOT	Above reference — either deliberate drive, or drift	trim input -n

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

HOT is not an error here. Once you are staged, driving into HOT is how you use this compressor — the readout is telling you how far past reference you have gone, which is exactly what you want to know. Set it, note the figure, and keep it consistent across takes.

ATTACK · RELEASE

ATTACK · 0.02 – 0.8 ms

How fast the FET clamps once the signal passes the fixed threshold. Note the range: the *slowest* setting here is under a millisecond. There is no slow attack on this plugin, because there was none on the circuit.

Fully fast catches the front edge of a transient and removes it — a snare loses its crack and becomes body. Backing off toward 0.8 ms lets the initial hit pass before the clamp lands, which is how you keep the stick and still compress everything behind it.

RELEASE · 50 – 1100 ms

How quickly gain returns. Short release re-opens between hits, which pulls up everything in the gaps — that is where density, room and pump come from. Long release stays down through a phrase and sits transparently behind the source.

Release interacts with tempo more than any other control here. A release that returns just as the next hit arrives locks to the groove; one that returns slightly late drags the feel.

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

SOURCE	ATTACK	RELEASE	WHY
Snare	0.3 – 0.8	100 – 200	Keeps the stick, compresses the shell behind it
Room mics	0.02 – 0.1	50 – 100	Flattens the hits and pulls the room up between them
Rock vocal	0.2 – 0.4	200 – 400	Holds the line forward without chewing consonants
Bass	0.1 – 0.3	400 – 800	Long release avoids modulating the fundamental
Electric gtr	0.3 – 0.6	200 – 500	Evens the part, keeps the pick attack

On low bass, a short release can track the waveform itself rather than the envelope, and you hear it as distortion rather than as compression. If a bass part sounds gritty at moderate reduction, lengthen the release before you reach for anything else.



RATIO · and all buttons in

Five caps rather than a knob: 4 · 8 · 12 · 20 · ALL. On the hardware these were physical switches, and only one could sensibly be down at a time — which is why holding several in at once became a famous accident rather than a feature.

RATIO	CHARACTER	REACH FOR IT ON
4	The gentlest setting, and the most even-handed	Vocals, bass, anything you want held but not squashed
8	Firmer grip, still musical	Drum busses, guitars, louder vocal performances
12	Audibly clamping; the colour comes forward	Snare, individual drums, aggressive vocal
20	Effectively limiting — a ceiling rather than a shape	Peak control, holding a part absolutely still
ALL	Every cap engaged at once — the misuse setting	Room mics, parallel drums, deliberate destruction

Ratio and INPUT GAIN work together, not separately. A higher ratio at the same drive gives more reduction and a harder knee; the same reduction can usually be reached from either direction, and the two routes sound different. If a setting is close but too aggressive, try dropping a ratio and adding drive before you change anything else.

ALL BUTTONS IN

Engaging every cap simultaneously puts the sidechain into a state the circuit was not designed for. The knee distorts, the ratio becomes non-monotonic, the attack and release curves shift, and the whole thing generates far more harmonic content than at any single ratio.

The result is unmistakable and completely uncontrolled: thick, distorted, forward, with the room and the decay hauled up level with the hits. It has been used on drums on hit records for fifty years and it is still the fastest route to a big, obviously-compressed sound.

ALL is not a louder version of 20. It is a different behaviour — do not treat it as the end of the ratio ladder. Reach for it when you want the effect, and use MIX to dose it.

HELP · ! and ?

Two small caps at the top right of the meter row, above OUTPUT. ! latches. Press it and the panel enters help mode: hover any control and a card rises with a one-line explanation of what it does. Press again to switch it off. It changes nothing about the audio.

? opens this manual.

Reading REDUCTION - OUTPUT GAIN

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

REDUCTION **-7.0** dB DEEP *ease GR*

STATE	REDUCTION	WHAT IT MEANS	TYPICAL USE
KISS	0 to -3 dB	Barely engaging — peaks only	Peak safety with the FET colour but little movement
RIDE	-3 to -7 dB	The classic working range — the panel calls it <i>classic 3-6</i>	Vocals, bass, guitars, drum busses
DEEP	-7 to -12 dB	Audibly compressing — you hear it working	Snare, aggressive vocal, when it is part of the sound
CRUSH	-12 dB +	Heavy reduction, the effect is the point	Room mics, ALL buttons in, parallel blends

The prescription beside the state tells you what to do to move between them. Because there is no threshold, *ease GR* means reduce INPUT GAIN or drop a ratio — those are the only two ways down.

OUTPUT GAIN

Makes up the level the compressor removed. 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 +8.2

-8.2 VU

HOT

trim output -7.5

+7.5 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 — and on a compressor that adds this much colour, that check matters more, not less.





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.

This setting earns its keep harder on FET Comp than on any other ZEROPINK compressor. A microsecond-fast clamp and an ALL-buttons ratio generate a great deal of harmonic content — which is the point — and every one of those harmonics is a candidate for folding.

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 — a bright kit through ALL buttons in is exactly that case. 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 — which is exactly what a FET stage is prized for. 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, and on a FET design it matters acutely:

Detection. With an attack measured in tens of microseconds, the whole point is to catch the front edge of a transient. A sidechain running at base rate can miss a peak that occurs between two samples — so the fastest compressor in the suite responds to the wrong event, and the crack you were trying to clamp passes through untouched.

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. The faster the clamp, the more abrupt that change: FET attack times are precisely the condition that generates the most.

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.

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 — and driving hard is how this plugin is used. Arriving at 0 VU before you add drive means the drive you add is deliberate and known, rather than inherited from 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 — and the FET colour you wanted arrives without the grit you did not.

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.


This is the control that makes the extreme settings usable. ALL buttons in at 100% mix is often too much of everything; the same setting at 30–40% keeps the density and the room lift while the dry signal restores the transients you just destroyed. On drums this pairing is the reason the plugin exists.

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

–18 is the default and the right choice for program material. The lower settings are for percussive instruments — and since this plugin lives on drums, you will use them. A snare or kick has a very high crest factor: the peak sits 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.

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:

1

Stage in. INPUT GAIN until the INPUT needle sits at 0 VU and STAGE IN stops prescribing. This is the reference you will drive from.

2

Pick a ratio. Start at 4 for anything melodic, 8 or 12 for drums. Choose the behaviour first; find the amount afterwards.

3

Drive it. Push INPUT GAIN past reference until REDUCTION reads RIDE — the panel's *classic 3-6* — or further if you want to hear it.

4


Shape with ATTACK. Decide whether you want the transient kept or removed, then set release to the tempo of the part.

5

Match the output. OUTPUT GAIN until STAGE OUT reads IN/OUT matched, then bypass and compare. Now the comparison is honest.

6

Dose with MIX. If the compression is doing what you want but has taken the life with it, pull the mix back rather than backing off the drive — you keep the colour and get the transients returned.



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

FET 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.

There is no threshold control. Correct — there was none on the circuit. Use INPUT GAIN to reach the fixed threshold and OUTPUT GAIN to make the level back up.

It sounds distorted. Some of that is the design. If it is more than you want: lower INPUT GAIN, drop a ratio, lengthen RELEASE, or raise OS RENDER — in that order.

My bass sounds gritty. A short release on low-frequency material tracks the waveform rather than the envelope. Lengthen RELEASE toward 800 ms.

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	FET, fixed threshold. Ratio 4 · 8 · 12 · 20 · all-buttons-in
Attack	0.02 – 0.8 ms
Release	50 – 1100 ms
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.