

OPTO Comp

OPTICAL COMPRESSOR · AU + VST3 · macOS

ZEROPINK
USER MANUAL



Three true-VU meters, three controls, one guidance strip. Everything else is footer.

We all want the perfect mix.

Before anything else: get to ZERO

01

OPTO 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.

So the first move is always the same:

1. Play the loudest section of the part.
2. Bring the INPUT needle to around 0 VU, drifting a little either side.
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 rather than each one compensating in turn.

INPUT GAIN on the plugin — for trimming from there, and for driving the optical element harder once the staging is right.

3. Only then reach for PEAK REDUCTION.

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.

Do this once, properly, and the compressor's controls will behave the way their labels suggest. Skip it and you will spend the session fighting the ratio.



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

What OPTO Comp is

02

A digital model of an optical compressor. In the hardware it emulates, gain reduction is done by a light source and a photocell: signal drives the light, the photocell's resistance changes in response, and gain falls. The photocell is slow and slightly imperfect, and those imperfections are the sound.

OPTO Comp reproduces that behaviour in software rather than reproducing any one manufacturer's unit. What matters is not the circuit diagram but what the mechanism does to audio — and it does things that a threshold-and-ratio design cannot express, which is why the controls here are not the ones you may expect.

It is program-dependent. The photocell's recovery is not one speed. It releases quickly at first, then slows as it settles — a fast initial recovery followed by a long tail. Loud passages leave the cell “tired” and it recovers differently than after a quiet one. The compressor is therefore always responding to what came just before, not only to what is passing through it now.

It has no ratio control, and that is not a simplification. In the circuit being modelled, ratio is not a setting — it emerges from the photocell's own curve and the amount of reduction you ask for. Push harder and the effective ratio rises on its own. Adding a ratio knob would mean inventing a parameter the mechanism does not have.

GOOD AT

Vocals, bass, acoustic guitar, program material, drum busses — anything where you want the level held without hearing it being held. It is a smoothing tool. It rounds and settles.


NOT GOOD AT

Catching fast transients. The optical attack is unhurried by design — a snare's crack is gone before the cell responds. That is FET territory. Using OPTO Comp for it lowers everything *after* the transient, which usually makes it stick out more, not less.

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. They deliberately do not respond to instantaneous peaks, which is why they read as musical level rather than as digital peaks.



INPUT — level arriving at the compressor, after INPUT GAIN. This is the one you set first. Target 0 VU.



PEAK REDUCTION — how much gain is being removed, reading right to left from 0. The needle should move like a musician breathing, not twitching. If it is slamming and holding, you are asking too much of an optical circuit.



OUTPUT — level leaving the plugin, after OUTPUT GAIN. When the compressor is set well and output gain is matched, this needle should sit around 0 VU alongside the input needle.

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 material genuinely lives there. Sustained residence there means your gain staging has drifted. PEAK 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

Trim into the compressor. This sets how hard the signal hits the optical element, and on an optical compressor that is the primary expressive control — more input means more reduction and a steeper effective ratio, because both come from the same physical mechanism.


Set it to reach 0 VU on the INPUT meter, then use it as your drive control.

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

STATE	MEANS	PRESCRIPTION
LOW	Below reference — the compressor is being underfed	push input +n
ON TARGET	At reference. This is where you want to be before touching anything else	✓ staged
HOT	Above reference — the stage is running too high	trim input -n

The number beside the prescription is the exact correction, so there is nothing to calculate. Follow it and the state resolves.

Each stage prescribes against its own control: STAGE IN sends you to INPUT GAIN, STAGE OUT to OUTPUT GAIN, with the sign carrying the direction.




PEAK REDUCTION

The amount of compression. Higher values ask for more gain reduction.

This is a single control that governs both depth and, indirectly, character. There is no threshold and no ratio because the optical circuit does not separate them. Turn it until the PEAK REDUCTION needle moves the way you want to hear.



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



KISS means the compressor is barely touching the signal. It is the first rung of a four-state ladder, with each state fixed to an amount of gain reduction:

STATE	REDUCTION	WHAT IT MEANS	TYPICAL USE
KISS	0 to -3 dB	Barely engaging — occasional peaks only	Gentle safety, preserving dynamics almost entirely
RIDE	-3 to -7 dB	Holding level transparently, the way a hand on a fader would	The everyday setting for vocals, bass, acoustic guitar
DEEP	-7 to -12 dB	Audibly compressing — you hear it working	When the compression is part of the sound
CRUSH	-12 dB +	Heavy reduction, the effect is the point	Parallel blends, room mics, deliberate flattening

The prescription beside the state tells you what to do to move between them — *ease PR* to back off, or its counterpart to push further.

LIMIT · OUTPUT GAIN

LIMIT

Engages a limiting mode on the optical element, steepening the effective ratio toward brick-wall behaviour rather than the gentler compression curve.

Because ratio on an optical circuit emerges from the mechanism rather than a setting, LIMIT is how you ask for the steep end of that behaviour without a ratio control existing. Use it when you need a ceiling held rather than a level shaped.

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.



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.

The STAGE OUT readout prescribes in both directions — *trim output +n* when you are short of reference, *trim output -n* when you have overshoot — and confirms when you arrive:

STAGE OUT	LOW	trim output +8.2	-8.2 VU
STAGE OUT	HOT	trim output -7.5	+7.5 VU
STAGE OUT	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.



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

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

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.

The split exists so you never have to choose between a responsive session and a clean master. Work at 4× — 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. Returns diminish steeply past 16×, so treat 32× 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 16× 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

08

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 in a compressor:

Detection. A sidechain running at base rate can miss a transient that peaks between two samples, so the gain never responds to the loudest thing in the passage.

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. 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.

This is why oversampling changes how a compressor *feels*, not only how it measures — the envelope itself is better resolved.

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 4× 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, so arriving at 0 VU before you compress is not only a metering discipline — it keeps the process in the range where there is less to fold in the first place.

The fold, drawn

09

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 · 0 VU REF · SKIN

10

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. Try heavy reduction at 30–50% mix on drums.


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 — vocals, bass, guitars, busses, mixes.

The lower settings are for percussive instruments. A snare or a kick has a very high crest factor: the peak sits far above the average the VU is reading. At –18 the needle can sit politely near 0 while the actual peaks are already close to the ceiling. Dropping the reference to –24 or –27 pushes the whole scale down, giving those transients room above 0 VU where they belong, and letting the needle tell you something useful instead of understating what is happening.

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

11

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.

2

Set the depth. PEAK REDUCTION until the reduction needle moves musically on the loudest passages. Watch the needle's motion more than its value — you are looking for it to breathe with the part.

3

Listen for the tail. Because release is program-dependent, the compressor sounds different after a loud phrase than a quiet one. Let it run through a full section before judging.

4


Match the output. OUTPUT GAIN until STAGE OUT reads IN/OUT matched.

5

Bypass and compare. Now the comparison is honest.

6

Consider MIX. If the compression is doing what you want but taking life with it, pull the mix back to 60–80% rather than reducing the compression.



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

OPTO 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.

The meters barely move. Your session is running well below –18 dBFS reference, or the INPUT GAIN has not been set. Follow the STAGE IN prescription.

The compressor is not catching my transients. It is not meant to. See “What it is not good at”.

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
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)
Latency	Oversampling-dependent — see table. 0 samples with oversampling off
Sample rates	Validated 11.025 – 192 kHz. Logic project rates: 44.1 · 48 · 88.2 · 96 · 176.4 · 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.