The Clipper

Three controls. Precise peak shaping.



Interface reference: The Clipper 0.1.12. Settings shown are an example.

What it does

The Clipper reshapes waveform peaks with a continuously variable soft-to-hard clipping curve. Use it to trim fast transients or add deliberate harmonic density. It has three controls: Input Gain, Shape and Ceiling. Native formats are AU on Mac, VST3 and AAX Native; the DSP + Native edition also provides AAX DSP for compatible Pro Tools HDX systems.

Install through Lydtrold Installer

Download **Lydtrold Installer** for your platform from lydtrold.com/installer. Close your DAW, select the plug-in and the available formats, then click **Install / Update Selected**. Reopen your DAW and insert the plug-in. Use Lydtrold Installer for later updates too.

Your first pass

1. Insert The Clipper on a track or bus. Start with **Input Gain 0 dB** and **Ceiling 0 dB**.
2. Choose Shape: higher values for a sharper knee, lower values for a broader soft knee.
3. Raise Input Gain gradually while watching CLIPPED and listening to transients.
4. Compare at a similar playback level. If you need a controlled final true-peak ceiling, place a true-peak limiter after the clipper.

Inside this guide

Page 2: the three controls. Page 3: clipping, antialiasing and latency. Page 4: the meters and practical workflow.

The three controls

Input Gain | 0.00 to +24.00 dB; default 0.00 dB

Adds gain before the nonlinear stage. More gain pushes more of the waveform into clipping and raises the level of portions that remain below the knee. Input Gain is not the amount clipped and is not compensated automatically. For example, +6 dB approximately doubles the incoming amplitude, but a quiet passage may still show little or no clipping.

Shape | 0.0 to 100.0%; default 100.0%

Controls the width of a quadratic transition into clipping. **100%** is hard clipping: the ideal curve is linear until the ceiling, then flat. Lower settings widen the soft knee, gradually bending the waveform before the nominal ceiling. **0%** is the widest soft knee, not bypass or no clipping. This is a continuous curve control, not a wet/dry mix.

Why soft clipping can act below Ceiling

The knee width is w = 0.95 x (1 - Shape/100). Relative to ceiling amplitude C, the ideal positive curve starts bending at C x (1 - w) and reaches its flat region at C x (1 + w). At Shape 50%, w = 0.475: bending starts at about 0.525 x C. With Ceiling 0 dB, that is about -5.60 dBFS **after Input Gain**. The displayed clipping meter can therefore move even if peaks are below the nominal ceiling.

Ceiling | -12.00 to 0.00 dB; default 0.00 dB

Sets the amplitude scale and threshold of the oversampled clipping curve. Lowering it makes clipping begin at a lower signal level. It is not a simple output fader: it changes the signal's position relative to the knee, so it can change both level and harmonic character. There is no automatic makeup gain.

Ceiling and final output peaks

Ceiling is the clipping-stage reference, **not a guaranteed final sample-peak or true-peak maximum**. Reconstruction filters can overshoot after the nonlinear stage. The output meter shows sample level, while the central PEAK readout shows clipping amount. Neither is a final true-peak compliance meter.

Adjustment and bypass

Drag the knobs or type into their value fields. Double-click a knob to restore its default. Use the host's plug-in bypass to compare with the latency-aligned dry signal; The Clipper has no separate front-panel bypass knob. It also has no attack, release, mix or oversampling selector.

Inside the clipping engine

Signal flow

Input > Input Gain > FIR upsampling > antialiased clipping curve > reconstruction compensation > FIR downsampling > output
In stereo, the same settings apply to each channel's waveform independently. There is no common compressor-style gain envelope that turns both channels down when only one peaks.

The ideal transfer curve

Let g = 10^(Input Gain/20), C = 10^(Ceiling/20), a = abs(g x input / C), and w = 0.95 x (1 - Shape/100). For w > 0:
a <= 1 - w: f(a) = a
1 - w < a < 1 + w: f(a) = a - (a - (1 - w))^2 / (4w)
a >= 1 + w: f(a) = 1
Restore the input sign and multiply by C. At w = 0, use f(a) = min(a, 1). This defines the curve; the implemented antialiasing and filters also affect individual output samples.

What antialiasing does

Clipping creates harmonics. Harmonics beyond the available bandwidth can fold back as aliasing. The Clipper combines oversampling with first-order antiderivative antialiasing (ADAA): it integrates the nonlinear residual over the interval between consecutive oversampled input values. A small reconstruction correction compensates the linear-region droop. This reduces aliasing; it does not claim to eliminate all aliasing or remove the harmonics that make clipping audible.

Automatic oversampling

At 44.1/48 kHz, the engine uses three cascaded 2x stages for 8x processing. At 88.2/96 kHz it uses two stages for 4x processing. These are automatic, sample-rate-dependent choices. Native and AAX DSP use the shared clipping algorithm; the host still controls where processing runs.

Timing and phase

The engine has a fixed **74-sample latency**: approximately 1.68 ms at 44.1 kHz, 1.54 ms at 48 kHz, 0.84 ms at 88.2 kHz and 0.77 ms at 96 kHz. Host bypass uses the same delay. Linear-phase FIR filtering preserves phase relationships while introducing delay and possible pre/post-ringing around transients. Parameter changes are smoothed; there is no compressor-style audio attack or release envelope.

Meters and practical use

INPUT and OUTPUT

The side meters show left/right sample peaks before Input Gain and at the final output. Their scale runs from -60 to +6 dBFS and their visible envelopes decay for readability. In mono, the channel presentation follows the mono signal. These are sample-peak meters, not LUFS or reconstructed true-peak meters.

CLIPPING history | four seconds

The central graph shows measured attenuation inside the antialiased nonlinear stage. It compares that stage's actual output against the time-corresponding linear reference after Input Gain and oversampling. The display uses the largest reduction across channels and oversampled evaluations, collected into 20 ms frames. Downward movement means more clipping; the graph covers 0 to 18 dB of reduction.

CLIPPED and PEAK

CLIPPED is the maximum reduction in the latest complete 100 ms interval. **PEAK** is the greatest reduction in the visible four-second window. PEAK is **not output level or true peak**. Numerical readouts can exceed the graph's 18 dB range. Soft-knee attenuation counts as clipping even below the nominal ceiling.

When playback stops

The graph advances with analyzed audio, including silence. When the host stops providing audio frames, the last trace and PEAK can remain visible for inspection. The side meters settle and CLIPPED clears after inactivity. Bypassed audio reports no new clipping.

Drums, buses and a final master

On drums, start with a harder Shape and increase Input Gain only until unwanted spikes are reduced. On a bus, try a softer Shape for more gradual bending, but listen for low-frequency distortion and changes in stereo balance. For a master, a small amount of clipping before a true-peak limiter can reduce the peaks the limiter must handle. These are starting approaches, not built-in presets.

If it gets louder, harsh or misleading

There is no auto gain: use a downstream trim for level-matched comparisons. Reduce Input Gain if transients lose impact. If the output exceeds Ceiling, remember that filtering can overshoot; use downstream true-peak control for a delivery ceiling. If CLIPPED differs from input-minus-output level, that is expected: it measures the nonlinear stage, not whole-signal loudness change.

Interface and support

The window remembers its size; double-click a free background area to reset it. Automate Input Gain, Shape and Ceiling in the host. Updates: lydtrold.com/installer. Support: hello@lydtrold.com. Include your plug-in version, host and operating system.