VITRIC AUDIO · MUSE HUB EDITION

Chamfer De-Esser

User guide · version 0.2.0

What it is

A de-esser that reshapes each ess on its own, tracks the performance level so the de-essing does not change when the singer does, and draws its actual threshold on the display. Click any ess and see exactly what was done to it. Precise mode reports 191 samples; Live reports zero. VST3, AU and AAX for Windows and macOS.

The threshold is relative and level-tracked, not a line you chase. Turn the singer up 18 dB or down 18 dB and the delivered depth moves by 0.5 dB or less at the command level, and 0.75 dB or less measured end to end, so a threshold set on the verse is still right in the chorus. Each ess is handled on its own and reshaped spectrally rather than ducked: identically constructed events come out identical, with residual below -120 dBFS on every style and mode. At matched delivered depth on the same takes, held-vowel body moves 0.00 dB on average, and a frozen 10 dB floor on ess over body is enforced where gain is applied, so Chamfer cannot be driven into a lisp.

The display is the instrument. The pale curve is your input, the lit curve is what leaves, and the shade between them inside the band is exactly the energy that came out; the input tap is latency-aligned to the output tap, so that shade cannot show reduction that did not happen. The band pane is where detection listens: drag its edges, drag the pane to move both and keep the width, scroll to widen or narrow around its centre, or type exact edge frequencies. The RIDE line across it is the detector's own tracked reference, K-weighted over about 250 ms, and esses visibly cross it when they trigger. Under the glass, the ess strip

records every detected event, its height the depth applied and its brightness the confidence. Click one and the card freezes that ess: its spectrum before and after, the removed energy shaded, the depth, the confidence and a marker at its centre frequency.

AMOUNT sets how far each ess is reshaped toward the voice around it, 0 to 150 percent, and RANGE is the hard cap on how far it may go, 0 to 24 dB. STYLE picks the shape of the cut: Smooth tilts the whole band, Split is the classic band cut, and Surgical reshapes each ess bin by bin toward the target. All three are gain shapes on one engine, so a style change glides and never relatches. SPEED trims the automatic ballistics by up to 50 percent either way, OUTPUT is a plus or minus 12 dB trim, and BYPASS is gain-matched and click-free. LISTEN auditions Band, Voice or Removed, latching on click and momentary on hold. LEARN plays the source, places the band where its esses actually live and calibrates the ride, and it refuses to guess when there is not enough sibilance to trust.

The drawer holds the detector: Relative or Absolute threshold, a plus or minus 12 dB offset, Precise latency at exactly 191 samples at 48 kHz or Live at zero, detection in Stereo, Mid Focus or Sides, channel link, sense, and individual attack and release trims. There are 24 factory starting points named for sources, from close pop leads to hi-hats and steel-string guitar, each measured on its matched real take before it shipped. A and B hold two full snapshots with the session. The worst style costs 0.43 percent of one core at 48 kHz.

In Live mode Surgical falls back to Split, because there is no lookahead, and the switch says so honestly. LEARN behaves the same way: given fewer than about 40 confidently sibilant frames it declines the material and changes nothing, rather than guessing at a calibration.

Built and used at a working recording studio. Support: support@vitricaudio.com

Key features

- A relative, level-tracked threshold: delivered depth moves 0.5 dB or less across ±18 dB
- Per-event spectral reshaping, so identically sung esses come out identical below -120 dBFS
- A frozen 10 dB floor on ess over body: zero lisp breaches across the measured battery
- The ride line on the glass is the detector's own threshold, K-weighted over 250 ms
- Every ess lands on a clickable strip, and its card freezes the pre and post spectra
- Smooth, Split and Surgical on one engine; Precise 191 samples at 48 kHz or Live at zero

Working the glass

The pale curve is your input, the lit curve is what leaves, and the shade between them inside the band is exactly the energy that came out. The band pane is where detection listens: drag an edge with the grip at the top of its hairline, drag the pane itself to move both edges and keep the width, or scroll over it to widen or narrow around its centre. Shift is a fine move, double-click resets the pane and double-clicking one edge resets just that edge. Click the pane and the LOW EDGE and HIGH EDGE fields open on the glass for exact frequencies. The RIDE line across the band is the tracked threshold, drawn live, and esses visibly cross it when they trigger. The ess strip under the glass records every detected event; click one to freeze its card and click the card to dismiss it.

The deck and the drawer

AMOUNT, RANGE, the STYLE switch, SPEED, the three LISTEN chips and OUTPUT sit along the deck, with LEARN, DRAWER, BYPASS and OPTIONS stacked at the right. Every knob takes the same grammar: drag to adjust, Shift to drag ten times finer, the wheel to step on that control's grid, double-click or Cmd-click to reset, arrow keys to nudge, and click a value to type an exact one. DRAWER opens the detector over the display: threshold Relative or Absolute, offset, latency Precise or Live, detection Stereo, Mid or Sides, plus link, sense and the two envelope trims. OPTIONS carries the tooltip switch and the keyboard shortcut sheet.

Typical uses

- **Lead vocals.** Set it on the verse and forget the chorus. The relative threshold tracks the performance level, so the delivered depth holds across an 18 dB swing either way.
- **Backing vocals and stacks.** Per-event reshaping keeps every layer's vowels intact while the esses come down, with Split and Smooth for stacks that want less surgery.
- **Dialogue and voice-over.** The click, breath and bright-vowel traps are standing tests, so plosives and breaths are not mistaken for esses; detector precision on the labelled battery is 1.000 at defaults.
- **Hats, overheads and acoustic guitar.** Starting points named for sources beyond the voice, each measured over its matched real take before it shipped: real reduction, zero lisp breaches, no dulling.

Full specification

- Relative, level-tracked threshold: a K-weighted RMS over about 250 ms, drawn on the display as the RIDE line
- Delivered depth moves 0.5 dB or less at the command level across a plus or minus 18 dB program sweep
- Measured end to end on rendered audio, delivered depth moves 0.75 dB or less across the same sweep
- Absolute threshold mode with its own ABS THR control, for chains that need a fixed floor
- Threshold offset, plus or minus 12 dB, for sources that want a different operating point
- Per-event spectral reshaping: identically constructed events come out identical, residual below -120 dBFS
- Frozen 10 dB floor on ess over body, enforced where gain is applied
- Zero lisp-floor breaches across the 8-take battery and all 23 preset audition renders
- Held-vowel body moves 0.00 dB on average at matched delivered depth, worst case 1.3 dB
- No chatter: at most 4 gain flips per cycle allowed, measured median episode 46 ms
- Click, breath and bright-vowel traps as standing tests; detector precision 1.000 at defaults
- AMOUNT, the target: 0 to 150 percent of the way toward the voice around each ess
- RANGE, the hard cap: 0 to 24 dB, whatever Amount asks for
- STYLE: Smooth (band tilt), Split (classic band cut), Surgical (bin by bin toward the target)
- Three styles on one engine, so a style change glides and never relatches or clicks
- SPEED trims the automatic attack and release together, plus or minus 50 percent
- Individual ATTACK and RELEASE trims in the drawer, on top of Speed
- OUTPUT trim, plus or minus 12 dB; BYPASS is gain-matched and click-free

- LISTEN: Band, Voice or Removed; click latches, press and hold is momentary, keys 1, 2 and 3
- LEARN places the band where the source's esses live and calibrates the ride
- LEARN declines material with fewer than about 40 confidently sibilant frames and changes nothing
- Band pane: drag either edge, drag the pane to move both and keep the width, scroll to widen or narrow
- Typed LOW EDGE and HIGH EDGE frequencies, each bounded by its neighbouring edge
- Guard band around the detection pane that absorbs the edges of each edit
- Display input tap latency-aligned to the output tap, so the removed-energy shade cannot overstate
- Ess strip: every detected event, height the depth applied, brightness the detector's confidence
- Inspect card: that ess before and after at its deepest hop, removed energy shaded, centroid marked
- DETECT: Stereo, Mid Focus or Sides, with a channel LINK control and a SENSE operating point
- LATENCY: Precise reports and measures exactly 191 samples at 48 kHz; Live reports and measures zero
- In Live mode Surgical falls back to Split, and the switch says so
- A and B snapshots that persist with the session, with a copy across
- 24 factory starting points named for sources, each measured on its matched real take
- Presets set the sound and never touch Listen, Bypass or the output trim
- Keyboard: drag, Shift for ten times finer, wheel steps on the control's grid, double-click to reset, arrows to nudge
- Shortcuts: 1, 2 and 3 latch the Listen modes, D the drawer, L Learn, B bypass, Esc closes, ? opens the sheet
- 0.43 percent of one core on the worst style, 46.6 microseconds per 512-sample block at 48 kHz
- Stereo and mono, 44.1 to 192 kHz
- VST3, AU and AAX on Windows 10+ and macOS, Intel and Apple Silicon

Installation on Muse Hub

Purchase, installation, updates, refunds and removal for this Muse Hub edition are all handled by the Muse Hub app. The plugin opens and runs directly: there is nothing to unlock inside it, and the status chip reads MUSE HUB. On Windows the plugin installs as VST3 and AAX. On macOS the installer places VST3, AU and AAX builds in the standard system plugin folders. If the plugin does not appear in a host, rescan plugins, confirm the host is 64-bit, and in Audacity enable it in the Plugin Manager.

Requirements: Windows 10 or later (64-bit) or macOS 11 or later, Intel or Apple Silicon. Settings and presets are the same as the edition sold at vitricaudio.com, so projects open in either edition.

Frequently asked questions

Does the de-essing change when the singer gets louder?

No. The threshold is relative to the tracked program level. Across a plus or minus 18 dB sweep, delivered depth moves by 0.5 dB or less at the command level and 0.75 dB or less measured end to end.

Can it be driven into a lisp?

No. A frozen 10 dB floor on ess over body is enforced where gain is applied, whatever the settings ask for. Zero breaches were measured across the 8-take battery and all 23 preset audition renders.

What latency does it add?

Precise mode: exactly 191 samples at 48 kHz, about 4 ms. Live mode: zero. Both are reported to the host and measured. Live has no lookahead, so Surgical falls back to Split there and the switch says so.

What does the ride line actually show?

The detector's own tracked reference, plotted at its level on the spectrum's scale. It is not a drawing of the threshold, it is the threshold. The display's input tap is latency-aligned to the output tap, so the removed-energy shade cannot show reduction that did not happen.

What does LEARN do?

Play the track for a few seconds and press LEARN. Chamfer places the band where this source's esses actually live and calibrates the ride to the voice. If there was not enough sibilance to trust, roughly 40 confidently sibilant frames, it says so and changes nothing.

Does it work on sources other than vocals?

Yes. The 24 starting points are named for sources, including hats, overheads and steel-string guitar, and every one was measured over its matched real take before it shipped.

What formats and DAWs does Chamfer De-Esser support?

VST3, AU and AAX on Windows 10+ and macOS, Intel and Apple Silicon. That covers Pro Tools, Logic Pro, Ableton Live, Cubase, Studio One, FL Studio, REAPER and every other major host.

Support: support@vitricaudio.com · Product page: https://vitricaudio.com/chamfer · Vitric Audio, Cleveland, Ohio. This guide describes Chamfer De-Esser 0.2.0 as distributed on Muse Hub.