TSL SERIES · DYNAMICS

LAP TSL-A34

TWO-CHANNEL JFET LIMITER / COMPRESSOR · V5.3 · USER MANUAL

> **How it works - and how it differs from a plugin compressor.** This is a feedback JFET limiter, and the controls mean what they meant on the unit. There is no threshold control: the threshold is fixed by the board's trimmers (about −6.5 dBFS at 10:1, falling to about −15.5 dBFS at 3:1). GAIN is the drive into the limiter (0 … +20 dB in front of the gain cell), not make-up gain; OUTPUT adds 0 … +6 dB after the board. RATIO moves threshold and slope together - towards 3:1 the unit starts earlier and more softly, towards 10:1 later and harder - and the knee is hard. ATTACK and RELEASE read the unit's scale; the actual gain change is slower. The first transient after silence passes at full GAIN until the envelope has charged, and at low frequencies the envelope ripple modulates the gain within the cycle. LIMIT lights from the threshold, GAIN lights when the input stage is overdriven. BYPASS is a relay bypass: the board keeps running.

LAP TSL-A34 is a two-channel JFET compressor / limiter built as a component-level circuit model. The plugin does not run a "compressor algorithm": the whole audio board (input followers, gain stage, the JFET-controlled stage, the precision rectifier, the attack / release network, the RATIO divider and the balanced output driver) is solved as a circuit in real time, component by component - including the operational amplifiers with their finite bandwidth, slew rate and input offsets, and the tolerance of every part. Each channel is its own board with its own set of part values, and the two boards are aligned to each other like a matched stereo pair - the channels compress alike and the stereo image stays in place. After the board sits a model of a real audio output transformer with a selectable drive (SATURATION).

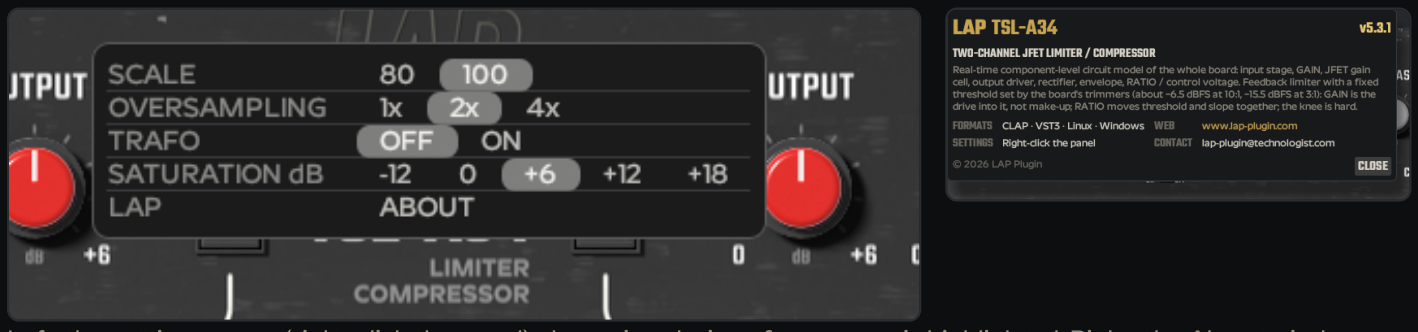
The two channels are two independent boards. **LINK** (the ON / OFF switch in the middle of the panel) makes channel 2 follow channel 1 for stereo work.

> **A mastering plugin.** TSL-A34 - like every plugin of the TSL series - solves the circuit in real time, and one stereo instance takes a sizeable part of a CPU core (about 19 % of real time at the default 2×, 48 kHz on the development machine). It is made for the master bus, for stems and for a few key tracks, not as a channel compressor on every track of a session: a handful of instances is fine, a dozen on busy tracks will run the CPU out. Use it where it counts.



1 Working the panel

Action	Result
drag a knob	vertical drag changes the value; hold `Shift` for fine control (1/10 speed)
hover a knob	the exact value is printed below the knob scale while the pointer is over the knob
double-click a knob	resets to default
click BYPASS	toggles the channel bypass; a pressed button with the red bar lit is bypassed. The switch cross-fades over 5 ms and the dry path is delay-matched, so it is click-free and usable while playing. The board keeps working while bypassed, as in the unit - the detector goes on following the programme, so the gain reduction is already in place when the bypass is released. For the same reason BYPASS does not lower the processor load.
click ON / OFF	LINK, see the chapter on LINK
click the LAP logo	opens a small About window with the plugin version, formats and links; hovering the logo shows the version
right-click the panel	the settings menu, see below
preset bar	above the panel, see Presets



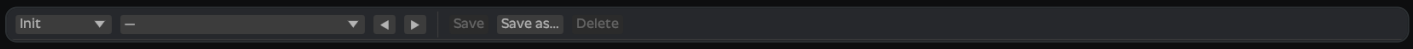
Left: the settings menu (right-click the panel); the active choice of every row is highlighted. Right: the About window (click the LAP logo) with the version, formats and contact.

Menu row	Choices	What it does
SCALE	80 · 100	window size in percent: 1052 × 176, **1315 × 220** (default), plus the 30 px preset bar
OVERSAMPLING	1x · 2x · 4x	see Oversampling and CPU; default **2x**
TRAFO	OFF · ON	the output transformer, see Sound and character; default **OFF**
SATURATION dB	−12 · 0 · +6 · +12 · +18	drive of the output transformer, see Sound and character; default **+6**. With LINK off the menu has one row per channel (SATURATION 1 / 2).
LAP	ABOUT	the About window

A click on SCALE, OVERSAMPLING or ABOUT closes the menu; TRAFO and SATURATION keep it open, so the settings can be compared one after another while the music plays. A click anywhere outside closes it. These are system settings - nothing on the panel shows them, the menu is the place to read them.

The panel follows the layout of a classic two-channel rack limiter and shares its frame with the other plugins of the TSL series: the control names are printed above the knobs, the range of each knob at the ends of its scale below it, and the two indicator LEDs sit in the top row - GAIN above the GAIN knob, LIMIT above the BYPASS button. The knobs stand in the order GAIN, ATTACK, RATIO, RELEASE, OUTPUT and turn through 270°; BYPASS follows OUTPUT next to the centre block with the LAP logo, the name and the LINK switch. Channel 2 is laid out as a mirror image of channel 1, so the OUTPUT knobs of both channels face the centre and GAIN sits at the outer edge. The whole panel - front plate, screen print, knobs, switches and LEDs - is a 3D model rendered with ray-traced light, at three times the resolution of the default window so it stays sharp when enlarged or on high-DPI screens.

2 Presets



Element	What it does
category ▾	choosing a category loads its first preset at once. Init is the default state of both channels (GAIN +3 dB, ATTACK 0.10 ms, RATIO 4:1, RELEASE 0.64 s, OUTPUT +3 dB, LINK on, output transformer off, SATURATION +6 dB, Oversampling 2x); User holds presets saved without a category; every other category is a folder of your own.
preset ▾	the presets of the chosen category; user presets are marked with a dot. A * after the name means a control or a menu setting was changed since the preset was loaded or saved; "-" means no preset is selected. The selected preset is saved with the project.
◀ ▶	step through all presets - quick A/B between saved settings
Save	overwrites the selected user preset with the current settings

Save as…	name and category (type a new one or pick from ▾); OK saves, Cancel returns
Delete	deletes the selected user preset - click twice (“Really delete?”)

A preset stores both channels, LINK and everything in the right-click menu: Oversampling, the output transformer switch and SATURATION. It does **not** change the BYPASS buttons and the window size (SCALE). Loading a preset moves the controls like a hand on the knobs: the audio keeps running without a drop-out, and the host records the change as automation if you are writing it. A preset saved by an older version still loads; controls it does not know start from their default (presets from before v5.3.1 load at 2x Oversampling, presets from v4.0 load with SATURATION 0 dB).

Presets are plain files: Windows `%APPDATA%\LAP\TSL-A34\presets\`, Linux `~/.local/share/LAP/TSL-A34/presets/` - one JSON file per preset, one sub-folder per category. They can be copied between machines and systems.

3 Controls (per channel)



Control	Range	Notes
GAIN	0 … +20 dB	gain of the input amplifier, before the controlled stage - it drives the signal into the threshold. 0 dB fully left; default +3 dB.
ATTACK	0.1 … 10 ms	attack time of the detector. The knob has a logarithmic law: every tenth of the travel lengthens the attack by the same factor. Shortest fully left - that is the default, 0.10 ms.
RATIO	3:1 … 10:1	on this circuit RATIO moves the **threshold** as much as the slope, see Levels, threshold and ratio; default 4:1. The knob can be turned and automated while playing: the plugin carries the state of the detector over to the new setting, so the level does not jump or dip when RATIO moves.
RELEASE	0.05 … 2 s	release time of the detector, linear law. Shortest fully left; default 0.64 s.

OUTPUT	0 … +6 dB	make-up gain after the board, ahead of the output transformer; 0 dB fully left is the unity-gain setting of the board; default +3 dB
BYPASS	ON / OFF	the board's own bypass: the clean signal from the input stage goes straight to the output, delay-compensated; OUTPUT and the output transformer do not act on the bypassed signal

> **ATTACK and RELEASE read the scale of the unit.** The figures at the ends of the scales - 0.1 … 10 ms and 0.05 … 2 s - are the time constants the unit is specified with at the end stops of its two pots, and the hover read-out follows the same scale. How long a gain change takes in practice depends on how deep it is and on RATIO: measured on the gain reduction with a 20 dB step at 4:1, the network needs about 5 ms (ATTACK min) … 120 ms (max) to reach 90 %, and about 0.2 s (RELEASE min) … 2.5 s (max) to recover to within 1 dB. Short settings are fast enough for limiting, long ones are programme-friendly.

4 Levels, threshold and ratio

Each channel's threshold is set by its internal trimmers, the way the board is aligned on the bench, so it sits at about **−6.5 dBFS with RATIO at 10:1** and moves down to about −15.5 dBFS at 3:1 - the same on both channels. Below threshold the board passes the signal at close to unity (about −1.1 dB; the 20 Hz – 20 kHz response of the coupling network is within ±1 dB); the output transformer takes about another 0.9 dB, so the whole plugin passes about −2 dB below the threshold with OUTPUT at 0 dB.

input dBFS	3:1	4:1	6:1	10:1
−24	−26.0	−26.0	−26.0	−26.0
−15	−18.6	−17.0	−17.0	−17.0
−12	−18.1	−16.4	−14.0	−14.0
−9	−17.5	−15.9	−13.4	−11.0
−6	−16.9	−15.4	−13.1	−9.6
0	−15.4	−14.3	−12.4	−9.3

Output level [dBFS] of channel 1 for a 1 kHz sine, GAIN 0 dB, OUTPUT 0 dB, output transformer on, SATURATION +6 dB, steady state. Channel 2 is within 0.1 dB.

The knee is hard and the slope above the threshold is steep; the printed 3:1 … 10:1 describes the whole curve from the threshold to +15 dB over it. Turning RATIO towards 3:1 is the way to get earlier, gentler action: the threshold drops and the first decibels are compressed softly. GAIN is the drive control: with the threshold fixed, GAIN decides how far into the limiter the programme goes; use OUTPUT to bring the level back.

Above the threshold the JFET adds a little harmonic distortion (about 1 % at 4 dB over at 10:1, 0.35 % at 3:1, almost all second harmonic) - that is the sound of the circuit, not an added effect.

5 Sound and character

Below the threshold the board passes the signal at unity and is very clean: the harmonic distortion is that of the amplifier chain - in the order of 0.001 % at 1 kHz, rising towards 0.02 % at 8 kHz as the loop gain of the amplifiers falls. Above the threshold the JFET takes over: about 1 % second harmonic at 1 kHz a few dB into limiting, and at low frequencies the ripple of the envelope modulates the gain within the cycle (several percent at 63 Hz when driven hard) - the classic sound of a fast feedback limiter.

Output transformer (TRAFO, default OFF). After the board the signal passes through a model of a real audio output transformer: the core with its magnetisation and hysteresis is solved sample by sample (at 4 × 44.1 / 48 kHz, i.e. 176.4 / 192 kHz whatever the session rate), driven from the balanced output of the board and loaded like a line input. It is the colour of the unit in this plugin - weight in the low end and low-order harmonics that grow with the bass content and the level, while mid and high frequencies stay clean. Switch TRAFO off for the plain line output of the board. The transformer adds 36 samples of latency at 44.1 / 48 kHz (30 at 88.2 / 96 kHz, none at 176.4 / 192 kHz, where it already runs at the session rate); switching it on or off while playing is click-free (the host is told about the new latency).

SATURATION sets how hard the transformer is driven: the signal is raised by the chosen amount in front of the transformer and lowered by the same amount after it, so the level stays the same and only the core works harder or softer. Measured on the transformer with a full-scale (0 dBFS) tone:

SATURATION	−12 dB	0 dB	+6 dB	+12 dB	+18 dB
THD at 40 Hz	0.38 %	0.34 %	0.34 %	0.46 %	43 %
THD at 100 Hz	0.14 %	0.15 %	0.14 %	0.14 %	0.16 %

The core saturates abruptly, like real iron: up to +12 dB the transformer adds only a trace of colour, at +18 dB the deep bass of loud passages goes into saturation (at 40 Hz the core clips from about +13 dB over full scale, at 90 Hz from about +20 dB). Lower levels and higher frequencies saturate much less. Changes of SATURATION are smoothed over 20 ms and can be made while playing.

First hit after a pause. This is a feedback limiter: with no signal the control voltage sits well below the JFET pinch-off, so after a few seconds of silence the gain element is fully open and the first transient passes at full GAIN until the detector has charged the envelope past the threshold - a couple of milliseconds at ATTACK minimum, close to a hundred milliseconds at ATTACK maximum, and longer towards 10:1 where the idle control voltage sits deepest. After that the loop settles and later hits are caught normally. It is part of the character of a feedback limiter, not a fault; if you need every first hit caught, keep some signal running into the unit, or use a look-ahead limiter after it.

Two boards, a matched pair. Every resistor and capacitor of each channel has its own value within the tolerance of its part series, together with the spread of the semiconductors and the amplifier offsets. The panel controls are the exception: the same knob position is the same setting on both channels. The trimmers of channel 1 are set by the alignment procedure of the board and channel 2 is then aligned to channel 1, so both channels have the same threshold and gain reduction (within about 0.1 dB), the same level (within about 0.2 dB) and closely matched timing, and the LIMIT LEDs move together. With LINK on the panel edits both, but they are not bit-identical - as two channels of real hardware never are.

6 LINK and the two channels



The **ON / OFF** slider in the middle of the panel is LINK (default ON). With LINK on, every knob, BYPASS and SATURATION of channel 1 is mirrored to channel 2 and the panel edits both at once; channel 2's own settings are copied from channel 1 the moment you switch LINK on, so nothing jumps when you switch it off again.

With LINK off the unit is dual mono: two boards, two sets of controls, no interaction between the channels, and the settings menu has a SATURATION row for each channel.

7 LIMIT and GAIN indicators

Each channel has two LEDs in the top row of the panel, above the LIMIT and GAIN captions. The red **LIMIT** LED shows the gain reduction: it comes on at 1 dB and grows brighter with the depth of the reduction, reaching full brightness at 8 dB and more; it fades over about a tenth of a second, so brief peaks stay visible. In silence it is dark at every RATIO setting. The green **GAIN** LED is the overload indicator of the input amplifier: it lights when the GAIN stage is driven into clipping. When it flickers on programme, turn GAIN down - the stage is clipping before the compressor. Both LEDs are dark on a bypassed channel and while the host is not processing audio (transport stopped, track disabled).

When the plugin is switched off and on in the host, or the transport is stopped and started, the detector continues from where it was, so the first hit is not let through louder. After a pause of more than a minute, and always in offline rendering, the board starts from rest like the unit after a pause (see Sound and character).

8 Oversampling and CPU

The default is **2×**: the response follows the circuit up to 20 kHz. At **2×** the audio stages of the board are solved at twice the host rate, which pushes the aliasing products of hard limiting further out. At **1×** the circuit is integrated with the trapezoidal rule at the host rate. The poles of the board that lie above half the sample rate would be warped by that integration into a roll-off near 20 kHz, so at 1× the plugin compensates for it with a small fixed equaliser - the response then follows the circuit to within 0.1 dB up to 0.375 × the sample rate (18 kHz at 48 kHz) and rolls off gently above (−0.8 dB at 19 kHz, −3 dB at 20 kHz at 48 kHz). 1× is the lightest setting for many instances, at about 3 % less CPU than 2×; **4×** is for 44.1 kHz sessions with very bright material. The choice is saved with the session. Switching it while playing is safe: the new model is built in the background and swapped in a few blocks later, without a drop-out, with a short cross-fade (about 21 ms) between the old and the new chain so the switch itself is click-free; the host is told about the new latency.

Sessions above 48 kHz. Oversampling is counted from 44.1 / 48 kHz, not from the session rate: the board and its detector always run at 1×, 2× or 4× of 44.1 / 48 kHz. In an 88.2 / 96 kHz session the input is therefore halved before the board and the output interpolated back (2× runs the board straight at the session rate, without any resampling); in a 176.4 / 192 kHz session it is quartered (4× runs straight). Each setting thus sounds the same and costs the same in every session, instead of a 192 kHz session costing four times as much. The resampling adds latency, which is reported to the host: at 96 kHz 61 / 0 / 30 samples for 1× / 2× / 4×, at 192 kHz 147 / 61 / 0 (plus 36 / 30 / 0 with the output transformer, which always runs at 176.4 / 192 kHz). With very small audio buffers at high session rates (64–128 samples at 192 kHz is only 0.3–0.7 ms) the two channel threads have little time to meet the deadline; 1× and 2× stay well inside it, 4× is best used with a buffer of 256 samples or more. Give the host real-time priority where the system allows it.

The audio path of the board is solved as four small circuit stages in signal order (input stage, GAIN amplifier, controlled stage with the JFET, output driver), each its own nodal system - the same physics as one big matrix, at a fraction of the cost. Since v5.3 the detector - rectifier, envelope and RATIO / control voltage - is a compact model built from the same components and their values, solved in closed form instead of as three more circuit stages. It sounds the same as the full detector (the difference stays more than 50 dB below the signal on music) and takes much less time, most of all on drums, where the detector used to work hardest on every hit.

Stereo, 48 kHz, output transformer on	1×	2×	4×
processing time / real time, programme at −18 dBFS	≈ 16 %	≈ 19 %	≈ 30 %
latency (samples)	36	66	72

Measured on the development machine for one instance (both channels); with the transformer off the load is about 1–3 % lower and the latency 0 / 30 / 36 samples. Oversampling multiplies only the audio stages (input, GAIN, JFET stage, output driver); the detector - rectifier, envelope and control voltage - always runs at the host rate, so 4× costs much less than 4×. The two channels are two independent models and are computed on two threads; a mono instance computes one channel only. The gain reduction, threshold and ratio are identical at 1×, 2× and 4×; only the top-octave response and aliasing differ.

Load follows the signal. A circuit model works harder when the circuit does: below the threshold each stage is solved in one pass, while a stage that is driven into clipping needs several. Programme material, full level included, stays well inside one core. Only material that clips the board continuously over the whole band - GAIN at +20 dB on a full-scale broadband signal, or inputs far above 0 dBFS - would need more than real time. For that case the plugin has an **overload guard**: in real time the solver works within a time allowance measured for your computer when the first instance starts. A single demanding sample (a steep edge) is still solved exactly from a small reserve; under continuous overload a stage that could not be finished holds its state for a sample - a little extra distortion where the board is clipping hard anyway, instead of a drop-out. Normal material never reaches the guard and is computed exactly as without it. Offline rendering (bounce, export) does not use the guard at all: there every sample is solved exactly, whatever it takes.

9 Signal flow

Stage	What happens
input	balanced input, input filter and input stage (unity gain)
GAIN	gain stage, 0 … +20 dB on the knob (its output drives the GAIN LED)
controlled stage	the gain cell: under the control of the detector the JFET takes up to about 17 dB out of the signal
OUTPUT · BYPASS	output stage at its unity-gain setting, BYPASS switch, balanced output driver
detector	from the output: full-wave rectifier, ATTACK charges and RELEASE discharges the envelope
RATIO / CV	the control amplifier sums the envelope with the RATIO offset and drives the JFET
after the board	1×: integration-warping compensation · OUTPUT make-up gain 0 … +6 dB · output transformer with SATURATION (solved at 176.4 / 192 kHz, can be switched off)