PHIL SPEISER

THE_STRIP 3

USER MANUAL

A complete channel strip that knows what it's processing — and every instance in your session knows about every other one. Includes **THE_STRIP 3 MINI**, the zero-latency companion you put on everything else.

Updated October 18, 2026
The latest version is always at philspeiser.com/manual/the_strip-3

CONTENTS

1 What It Is

2 The Rack You're Replacing

3 Why You'll Love It

4 The Big Idea (read this once)

5 The Interface at a Glance

6 SMART SETUP — the strip that configures itself

7 Calibration — why the numbers mean something

8 The Rack

9 The Fifteen Modules — In Depth

10 Output & Gain Staging

11 The MIX View — one mixer for the whole session

12 UNMASK — decide who wins

13 THE MIX MAP — your session as a room

14 AUTO-MIX — the one-button pass

15 PHASE ALIGN — one press, every mic in time

16 THE_STRIP 3 MINI — the companion

17 Visual Feedback & Metering

18 How to Use It — Real-World Workflows

19 Pro Tips & Creative Techniques

20 Keyboard Shortcuts & Gestures

21 Quick Start

22 Specs

What It Is

THE_STRIP 3 is an end-to-end channel strip in a single plugin: input staging, low-cut, gate, six-band dynamic EQ, compressor, de-esser, sidechain duck, transient designer, multiband compressor, stereo imager, reverb, delay, modulation, analog warmth and a brickwall limiter — **fifteen modules fused into one engine**, reorderable by drag, each with a full-depth editor behind its card.

Three ideas make it different from a folder full of plugins:

1. **It listens.** Pick a source type — or let it decide — and hit **LISTEN**: SMART SETUP analyses the actual audio on the track, configures every stage, and loudness-matches the output. Eighty-one factory profiles from LEAD VOCAL to the INSTRUMENTAL beds, plus profiles you create from your own reference audio.
2. **It's calibrated.** The strip gain-stages itself to a **-12 dB operating reference** at every point in the chain, and once it's CALIBRATED, the threshold numbers mean exactly what they say — the compressor's default readout is literally "how many dB of gain reduction do you want".
3. **It's a mesh.** Every THE_STRIP 3 (and MINI) instance in your session finds the others automatically. Flip to the **MIX view** and you're looking at a cross-instance mixer: faders, pans, meters, solo, groups — for the whole session, from any window. On top of that mixer sit the three session-level tools: **UNMASK** (a pecking order — drag to decide who wins, and the losing tracks' dynamic EQs step back only while the winner plays), **DUCK** (any track pumps under the session's actual kick, no routing), and **AUTO-MIX** (one button that sets up, unmasks and balances the whole session — tracks first, then your buses, then your master).

THE_STRIP 3 MINI is the same engine narrowed to the essentials — input, low-cut, compressor, duck, and the network's EQ — with **zero latency by construction**. It's the citizen you put on forty tracks so the full strip on your bus has something to mix.

The Rack You're Replacing

Every one of these is a processor you'd otherwise buy, insert and route separately:

Module	What you'd normally reach for
INPUT	A gain/trim utility with polarity flip and L/R swap
LOW-CUT	A resonant high-pass — a true Butterworth at 6/12/24/48 dB/oct, -3 dB at the corner at *every* slope
GATE	A transient-smart gate/expander with hysteresis, band-limited detection, external sidechain and lookahead
EQ	A six-band **dynamic** EQ with per-band Mid/Side/L/R targeting — and, on the network, a sidechain from *another track's* spectrum
COMP	A compressor with Clean/Opto/Punch envelopes, parallel mix, M/S modes, external sidechain — and a threshold that reads in dB of reduction
DE-ESS	A split-band or wideband de-esser with M/S targeting
DUCK	A sidechain ducker that pumps on the host grid, on a real sidechain, or under the session's kick — with a drawn recovery curve and frequency focus
TRANSIENTS	A transient designer with real attack/sustain windows, contrast, band focus and per-stream drive
MB-COMP	A three-band multiband compressor, full per-band control
STEREO	An M/S imager with bass-mono, three-band width, monoize, balance and a dynamic width governor
REVERB	An algorithmic reverb with per-band decay — and a built-in **wet ducker** so the tail steps back while the dry plays
DELAY	A tempo-synced stereo delay with ping-pong, filtered feedback, saturation, cross-feed, wow — and its own wet ducker
MOD	Chorus, flanger, phaser, tremolo and auto-pan over one LFO, free-running or locked to the host grid
WARMTH	Six analog saturation models with bias, tone, tape hysteresis and wow, independently oversampled
LIMITER	A lookahead brickwall limiter with three characters, clip safety, oversampling and an optional true-peak stage

And the honest version of the pitch isn't that it's cheaper than buying them. It's that owning all of them still wouldn't get you SMART SETUP writing the whole chain from a listen, thresholds that read in real dB because the strip staged itself, or a mixer where the vocal's EQ ducks the guitar bus *only while the vocal plays* — across plugin instances, with no routing.

Why You'll Love It

- **One LISTEN pass to a working chain.** Pick VOCAL, play the loudest phrase, and the strip writes the low-cut, the de-esser, the EQ (including dynamic notches on the resonances it actually found), the compressor, the character and the output level. Everything stays hand-editable.
- **Thresholds that tell the truth.** Calibrated, the compressor's default mode reads in **dB of gain reduction**: set -4, get 4 dB on the loudest material, at any ratio. No other knob-guessing loop.
- **Drag-to-reorder rack.** Thirteen of the fifteen stages re-sequence by drag; INPUT stays first (it's the reference) and LIMITER stays last (it's the guarantee).
- **Delta-solo any stage.** The **S** chip plays exactly what that module contributes — the difference, not the output. The fastest way to hear what the de-esser is actually removing.
- **A mixer for the whole session, in the plugin.** Volume, pan, mute, solo, groups, loudness, true peak and correlation for every instance — plus HYBRID, which docks that mixer as a rail above the open channel.
- **UNMASK.** Drag tracks into a pecking order; everything below steps back for what's above it, with dynamic EQ that only moves while the winner plays. Kick and bass never get carved — that pocket goes to the duck instead.
- **AUTO-MIX.** One button: learns every strip that needs it, unmasks the tracks, arms the kick ducks, settles the levels tier by tier — tracks, then buses, then master — and pulls the sum to target. One Undo takes back the entire run, on every instance.
- **Honest A/B everywhere.** Bypass is level-matched by default (MATCH), every module's output trim keeps per-stage A/Bs fair, and the loudness readouts are always post-everything.
- **A companion built for scale.** THE_STRIP 3 MINI: five stages, one hero pad, zero PDC — and a full mesh citizen: it gets carved, ducked, metered, drilled into and learned remotely like a full strip.
- **Multilingual tooltips (EN/DE/FR/ES/IT), full undo/redo, UI zoom, in-app docs and an AI assistant.**

The Big Idea (read this once)

Three words explain the whole product.

THE RACK — fifteen modules, one engine, one signal path you can reorder. A card is a module; click it and the full editor opens below. The gaps *between* cards are real: each carries that stage's live meter, its measured in→out delta, and an output trim — so you can see and correct what every stage does to your level, right where it does it.

THE REFERENCE — the strip stages itself to **-12 dB** at every module boundary. Switch **CALIBRATE** on (or run SMART SETUP, which does it for you) and the strip measures the input, trims it to the reference, then keeps every stage handing the next one the level it received. Once the chip says **CALIBRATED**, threshold knobs across the strip read against that reference — the numbers mean what they say. Nothing here is a continuous servo: the strip only ever measures in short windows after *you* change something, then holds. **The strip never changes your audio on its own.**

THE MESH — every instance of THE_STRIP 3 and THE_STRIP 3 MINI in the session connects automatically. The mesh is what makes the MIX view a real mixer, UNMASK a real arrangement of who wins, DUCK a real "pump under the session's kick", and AUTO-MIX a real one-button mix pass. No routing, no sidechain cables, no setup.

Everything else in this manual is detail on top of those three words.

The Interface at a Glance

One window, left to right and top to bottom:

- **Header row** — on the left, the **setup stamp** once the strip has been set up: **SET UP AS LEAD VOCAL** for a profile you picked, **AUTO HEARD LEAD VOCAL** for one AUTO recognised while listening. It stays until the next LISTEN, comes back with the session, and clicking it opens SMART SETUP. Then the **SMART** chip (opens SMART SETUP), the **calibration chip** (NOT CALIBRATED / CALIBRATING… / WAITING FOR AUDIO… / CALIBRATED — click it for the LEVEL CALIBRATION popover), and the **[STRIP | MIX | HYBRID]** view toggle (HYBRID appears once a second instance joins the mesh).
- **Top bar** — tooltip readout, undo/redo, AI assistant, documentation, settings, bypass. The corner grip zooms the whole UI.
- **The rack** — the module cards in signal order, scrolling horizontally. Enabled cards show live visualisations and two quick knobs; the selected card blooms wide; a disabled card collapses to a slim rail. Between every pair of cards: that stage's meter, delta readout and output trim.
- **The advanced editor** — the bottom two-thirds. Click any card and its full editor opens here. On a short window it collapses automatically (and can be collapsed by hand); the rack takes the space.
- **The XY panel** — the right-hand pad: vertical is output gain (±12 dB), horizontal is pan. Padlocks pin either axis. Under the puck, the final output's stereo field breathes as a particle cloud.
- **OUTPUT rail** — the slim card after the pad (L/R meter over the word OUTPUT). Click it — or TAB past the last module — for the OUTPUT big view: spectrum, loudness suite, gain map and the master action rail (CALIBRATE / RE-MEASURE / DIM / SET / BYPASS / MATCH).
- **MIX view** — the cross-instance mixer (see its own chapter).

SMART SETUP — the strip that configures itself

Press **SMART** in the header (or key **S**). The panel asks one question: **"TELL THE STRIP WHAT THIS CHANNEL IS."**

The signal-type grid

A full-width **AUTO** row ("NOT SURE? THE STRIP DECIDES WHILE LISTENING") and then 81 factory profiles by category:

- **VOICE** — LEAD VOCAL, RAP VOCAL, LEAD VOCAL (MALE), LEAD VOCAL (FEMALE), BACKING VOCALS (MALE/FEMALE), **VOCAL DOUBLE** (a stacked take under a lead — sits lower, goes wide on the stage spread), VOCAL BUS (MALE/FEMALE), SPOKEN / PODCAST, VOCAL CHOPS (an instrument made of voice — no de-ess, wide, delayed), CHOIR
- **DRUMS** — **DRUM KIT** (a whole stereo kit recording or loop on one track — it ranks, carves and solos like a track; a track just named "Drums" or "Kit" lands here, not on a bus), DRUM BUS (plus ACOUSTIC / ELECTRONIC / HIPHOP), TOP LOOP (ELECTRONIC / HIPHOP), KICK DRUM (plus ACOUSTIC / ELECTRONIC), SNARE (plus ACOUSTIC / ELECTRONIC), HI-HAT (ACOUSTIC / ELECTRONIC), CLAP, PERCUSSION LOOP, PERCUSSION HIT, PERCUSSION BUS, OVERHEADS, TOMS, RIDE / CYMBALS, ROOM MICS (a far mic — PHASE ALIGN only flips its polarity, never drags it into time)
- **BASS** — BASS, ELECTRIC BASS, SYNTH BASS, 808, SUB BASS, BASS BUS
- **GUITAR & KEYS** — ELECTRIC GUITAR (STRUMMED / PICKED), ACOUSTIC GTR (STRUMMED / PICKED), NYLON GUITAR (PICKED), GUITAR BUS, ACOUSTIC PIANO, POP PIANO, ELECTRIC PIANO, ORGAN, MALLETS, KEYS BUS
- **SYNTH** — SYNTH, SYNTH CHORDS, SYNTH MELODY, SYNTH PLUCK, SYNTH PAD, SYNTH BUS (CHORDS), SYNTH GROUP (MELODY)
- **ORCHESTRAL** — ORCHESTRAL (CHORDS / MELODY), ORCHESTRAL BUS
- **BRASS & STRINGS** — SAXOPHONE, TRUMPET, HORN SECTION, VIOLIN, CELLO
- **FX** — FX HIT, FX BUS (an Ableton-style return track named "Return" lands here automatically)
- **2-TRACK** — **INSTRUMENTAL** (POP / HIP-HOP & TRAP / R&B & SOUL / ELECTRONIC & DANCE / BAND & ROCK / ACOUSTIC): the whole backing track a vocal sits on. A bed is a *track*, not a bus — it is ranked under the voice and carved where the voice sings — and its LISTEN measures how finished the material already is: a bought, limited beat is left alone (level and carves only, the ledger says "already mastered - voicing kept light"), a rough bounce gets the glue and tone a mix bus would. A track named "Beat", "Inst", "Instrumental" or "Playback" lands here by itself (a bare "Beat" is the HIP-HOP bed). See **AUTO-MIX → Two tracks**.
- **MIX** — GROUP BUS, MIX BUS, MASTER, MASTER LOUD
- **YOUR PROFILES** — profiles you made, plus **+ CREATE PROFILE** (right-click a user profile to rename or delete).

The panel opens **pre-selected from your DAW track name** — a strip on "Lead Vox" doesn't open on AUTO. It's a draft, never an override; DAW master-bus names like "Main" or "Stereo Out" resolve to MIX BUS (only a "loud" in the name lands on MASTER LOUD).

Search the grid. A search well sits under the title, in SMART SETUP and in the mixer's type picker alike. Click it and type — or just start typing, the panel is listening the moment it opens. "kick" leaves the three KICK DRUMs under DRUMS, "drums" the whole category: a query matches category names as well as profile names. RETURN takes the first result, the x at the right clears, and ESC unwinds one layer at a time — the search first, then the panel. There is no text field to fight your DAW over; keys type into the panel itself, and a leading space still reaches the transport, so play and stop work while the panel is up.

A word that names no profile still gets an answer. The search reads your word the way the mixer reads a track name, so "ukulele" offers ACOUSTIC GTR (PICKED) and (STRUMMED), "hats" both HI-HATs, "vox" the three LEAD VOCALs, "vocoder" VOCAL CHOPS and both BACKING VOCALS, "jazz" the acoustic bed and then the piano, sax, bass and ride it is made of. Genre words open with their bed, synth names (juno, moog, serum, supersaw) land on their synth families, and instrument words in German, French, Spanish, Italian and Portuguese (Gesang, batterie, bajo, cuerdas, coro) are understood too. Those fits lead as *CLOSEST FIT FOR "..."* when nothing matches by name, and follow the name matches as *ALSO FITS "..."* when something does — "kick" reads the three kicks, then 808 and SUB BASS. When there is truly nothing, the grid says so: *NOTHING CALLED "..." - ESC CLEARS THE SEARCH.*

LISTEN

- **LISTEN FOR** — capture window: **3 s / 6 s / 12 s** chips. The window counts *material*, not time: capture is signal-gated, so silence pauses it ("WAITING FOR SIGNAL") rather than diluting it. Play the loudest part.
- **MAKE IT WET** — a sticky per-strip switch: the next learn guarantees reverb and delay switched on — ducked and tempo-synced — even where the profile would stay dry.
- Press **LISTEN** (later: **RE-ANALYSE**). The header chip shows a progress ring; the panel shows the live "PEAK SO FAR". Then a LEVEL pass ("BALANCING OUTPUT - HANG ON") holds until the strip's gain windows actually settle.

What a learn writes: input trim seeded from the true stereo peak — and from **section memory**: the strip remembers the loudest section it has heard in the last minute of playback, so a LISTEN pressed during a quiet breakdown still sets the level the chorus needs (the report says so when it happens); a **safety low-cut** where the profile authors one (and cleared again by a re-learn as a type that doesn't, so a hi-hat re-learned as a kick loses its 200 Hz filter); the **strongest three** tonal EQ moves plus **dynamic notches on resonances it actually measured** (it always leaves at least three EQ bands free for the network); compression tiered SOFT/MEDIUM/HEAVY from the material's crest — aiming at a musical 3–4 dB of reduction; de-esser threshold from the real HF envelope; transients, width (with a measured mono/correlation correction), saturation; reverb and delay per profile — the delay always tempo-synced; on the 2-bus profiles a **limiter push** into the -1 dB wall — MIX BUS 1 dB, MASTER 2 dB, MASTER LOUD 4 dB of steady reduction on the loudest material, level-neutral end to end — which every other profile writes back to zero; and the output level matched to the reference. Every press lands a tasteful variation, not a fixed template — and every learn is **one undo step**.

What a learn switches OFF: **GATE**, **DUCK**, **MB-COMP** and **MOD**. No profile carries any of them, so a learn can never switch one ON — and it now switches them off when it applies a profile that does not use them, because the alternative was worse: an extreme multiband you had dialled and forgotten would sit there contradicting the card that says SET UP AS MIX BUS, and nothing anywhere said so. **Your settings are not touched** — only the engage flag moves, so switching one back on gives you

exactly the sound it had — and every bypass is named, in the setup note and in the AUTO-MIX report. DUCK is the one with an automatic hand: AUTO-MIX arms the factory kick duck on the bass-family tracks under the kick (or on the grid, if you answered FOUR TO THE FLOOR), so on those tracks a run resets the duck and then re-arms its own — AUTO-MIX decides the duck rather than inheriting a stale one. Outside a run, a duck you switched on yourself is still never overwritten. (A learn also arms MB-COMP's auto-makeup, the same as pressing AUTOGAIN, so a band you switch on afterwards is already gain-staged.)

If it can't hear, it tells you which problem you have, exactly:

- Transport stopped — "NO SIGNAL - PRESS PLAY" / "Nothing was playing, so nothing was changed."
- Rolling but silent — "NO AUDIO ON THIS TRACK - check the track is unmuted and routed here."
- Audible but under the analysis floor — "TOO QUIET TO ANALYSE - RAISE THE LEVEL, THEN LISTEN."

The result is a plain-language report ("SMART SETUP COMPLETE") with the real numbers — "I set the compressor to 2.3:1 around -24.1 dB", "I tamed 3 ringing resonances with dynamic notches" — in signal order, always ending: *"Not what you wanted? Undo takes the whole setup back."* After you dismiss it, the profile stays on screen: the header's **setup stamp** reads *SET UP AS …* (your pick) or *AUTO HEARD …* (what AUTO recognised), the AUTO row of the grid says *LAST LISTEN HEARD …*, and the mixer card's type chip shows the same profile in every host, whether or not the DAW hands the plugin a track name.

CREATE PROFILE — teach it your sound

Three steps, and a rail across the top says which one you are on:

1. **PICK AUDIO** — drop or choose 1–8 files (WAV, AIFF, FLAC, MP3, OGG) that already sound the way this channel should. They land in a list you can add to and remove from before anything happens; **ANALYZE** starts the measurement, and each file reports back as it is heard (a tick, or the reason it could not be used). The strip measures their tone, dynamics and stereo image.
2. **DESCRIBE** — the panel first shows you **WHAT IT HEARD**: the measured tone curve, the crest, the stereo width, in plain words with the numbers under them. Then say what the sound *is* in free text ("aggressive rap lead vocal") and the strip works out how to treat it — the compression style, the de-essing and the character your files cannot tell it. Needs internet; offline it goes by the sound alone, and **SKIP** does the same on purpose.
3. **NAME IT** — name it and save. The suggestion comes from your own words (or, if you skipped, from the file you brought). The profile lives on this computer and appears under YOUR PROFILES in **every instance and every session**.

A profile you build from bus or mix references **is** a bus: if the strip recognised your files as a MIX BUS, DRUM BUS or MASTER, a card set to that profile behaves like one in the mixer — group solo, out of the pecking order, learned after its tracks have settled — exactly as the built-in bus types do.

MANAGE — the profile library

Every profile you make lives in one library on this computer, and **MANAGE** opens it: the pill on the YOUR PROFILES heading in SMART SETUP, and the same pill in the mixer's type picker. Left, every profile in the order the pickers show them — drag the grip at the left edge of a row to reorder, click a row to open it, double-click to use it on this strip. Right, the profile itself.

What a profile is. Three editable fields and one measurement. **NAME** and **DESCRIPTION** are yours to change at any time; a name that is already taken is refused with a message and the field keeps what you typed. **SIGNAL TYPE** is what the strip understood your files to be — it decides the compression style, the de-essing and the character. **CHANGE** re-aims the profile at a different signal type without touching the audio again: the measurement is kept, only the interpretation moves, so a profile built on the wrong signal type is a two-click fix instead of a rebuild. That signal type reaches the mixer too: use a profile of your own and its card takes that type — so a profile you based on MIX BUS behaves as a bus everywhere (its role, its group and colour, its place in the pecking order, its solo scope, and which tier it learns in during an AUTO-MIX run), exactly as a factory pick would. Under it, **WHAT THE STRIP HEARD** — tone, dynamics and stereo image against how that kind of source normally sits — and **MADE FROM**, the files it was measured on and when.

The actions.

- **USE** — select this profile on the strip you opened the library from, ready for LISTEN.
- **DUPLICATE** — a copy with the same measurement and its own name. Re-aim the copy and the original stays as it was.
- **EXPORT** — write this one profile to a .stripprofiles file. That file is the whole profile: measurement, base, name, description.
- **DELETE** — asks once (SURE?), then moves the profile to a trash on disk. A **RESTORE** chip puts the last deleted one back, and it still works after the plugin has been closed and reopened. The trash keeps the twelve most recent deletions.
- The three icon chips beside NEW PROFILE: **IMPORT** a bundle, **EXPORT ALL** profiles to one bundle, and **SHOW FILE** — reveal the library on disk.

Sharing and backup. A .stripprofiles bundle holds one profile or the whole library. Drop one anywhere on the library to import it, or send it to another computer. Import merges: a profile you already have is updated in place, a different profile that happens to share a name is kept under a new name — nothing is overwritten. The library is one file, user_profiles.json in the THE_STRIP folder of your application support directory (SHOW FILE takes you there); copy it and your profiles survive a new machine. Every instance and every session on this computer reads the same library, so a profile made in one song is there in the next — and a change made in one plugin window shows up in the others within a second.

Identity, not position. A strip remembers the profile it was set up with by identity, not by its place in the list. Reorder the library, delete a neighbour, import twenty more — the card, the grid and the saved session still point at the profile you chose. Delete the profile itself and the strip says so instead of quietly picking whatever took its place.

Calibration — why the numbers mean something

The header chip is the strip's honesty meter:

- **NOT CALIBRATED** — your gain staging is your own; threshold numbers are approximate.
- **CALIBRATING...** — measuring right now; levels are still moving.
- **WAITING FOR AUDIO...** — a measurement is open with nothing to hear. Press play and it lands by itself.
- **CALIBRATED** — every stage sits on the reference. Your level will not jump as you work, and the threshold numbers mean exactly what they say.

Click the chip for the **LEVEL CALIBRATION** popover: **CALIBRATE** (the switch), **RE-MEASURE** ("listen again and re-trim" — use it when the *audio* changed: a new clip, another section, a re-tracked take), **SMART SETUP**, and **GAIN-STAGING TARGETS** (target -24...-3, default **-12 dBTP**; window 0.5–6 s, default 2.5 s; range 3–24 dB). While a measurement is starved, RE-MEASURE becomes **STOP WAITING** — keep the trims, stop listening.

How the ladder works

Calibration is three **measure-and-hold** windows, each ~1.5 s of actual audio — never a continuous servo:

1. **INPUT** — measures the raw input peak and sets the input trim so the strip starts at -12.
2. **STAGE** — armed whenever you *edit* a module, and scoped to that module: it measures the stage's own in→out change and writes the stage's OUT trim to cancel it, so every stage hands the next one the level it received. (DUCK is never trimmed — its dip is the point.)
3. **FINAL** — sets the **GAIN CORRECTION** (visible on the OUTPUT page) so the finished chain lands back on the reference — or back on your original input level, if you set **RESTORE INPUT VOL**.

Everything the ladder is about to rewrite wears a **breathing steel ring** first — the trim lane in the rack, the OUT knob in the editor, the header chip. You always see the measurement coming. Silence never lands a window; state restores never re-measure; and nothing moves unless you touched something.

Threshold semantics

On a calibrated strip, the threshold knobs of the gate, de-esser, multiband bands and dynamic EQ bands read **relative to the -12 operating peak**. The compressor goes one better — its **THRESHOLD MODE** switch:

- **REDUCTION** (default) — "-4" gives you **4 dB of gain reduction** on the loudest material, at any ratio. The strip works out where the threshold has to sit, compensating for the detector you chose.
- **THRESHOLD** — classic: "-4" puts the threshold 4 dB below the reference and the ratio decides the reduction.

Only two numbers in the strip stay absolute on purpose: the DUCK's SC THRESH (it keys a foreign signal) and the LIMITER's CEILING (it's an output law, not a music decision).

The Rack

- **Signal order** (default): INPUT → LOW-CUT → GATE → EQ → COMP → DE-ESS → DUCK → TRANSIENTS → MB-COMP → STEREO → REVERB → DELAY → MOD → WARMTH → LIMITER.
- **Reorder by drag** — pick up any card body and carry it; the other cards ease aside and a recessed well marks the drop slot. **INPUT is pinned first** (it's the trim, the meters and the reference tap) and **LIMITER is pinned last** (the ceiling guarantee); the thirteen stages between them go anywhere. EQ before COMP is a different record from COMP before EQ — the order is a sound.
- **Enable LED** — click to switch a stage on or off. Toggles are ~20 ms crossfades, so nothing clicks, and reverb/delay tails fade out instead of cutting. A stage that's off collapses to a slim rail — OFF reads as OFF — and its output trim returns to unity so a dead stage can't keep faking a level.
- **Delta-solo (S chip)** — on every stage except INPUT and MOD: hear **only what this stage contributes** — its output minus its input. On DELAY it's the wet only; on WARMTH it's the added colour only; on the LIMITER it's the clip, the push and the wall together (a disabled limiter's delta is honest silence). Delta solos are editor gestures: they never survive closing the window.
- **Quick knobs** — each card carries its two headline parameters (COMP: threshold + ratio; REVERB: mix + decay; …). Same gesture law as every knob in the product: drag (140 px = full range), **SHIFT** for fine, **double-click** to reset, **right-click** to type ("1.52k", "-6 dB", "30%" all parse).
- **Between the cards** — every gap shows the left stage's live meter (RMS body, peak cap, and a GR pip descending from the top on dynamics stages), its measured **in→out delta**, and a bipolar **output trim** (±12 dB; gold above unity, steel below). Right-click the trim lane to type a value. This is where "what did it do" and "what did I add" stay two different questions.
- **HOT warning** — if any point between modules touches 0 dBFS, the XY panel warns you. The fix is usually one trim lane to the left.
- **TAB** walks the chain in signal order (Shift-TAB back), lands on **OUTPUT** after the last module, and wraps.

If everything is off, the rack says so: *"NOTHING RUNNING YET — CLICK A STAGE TO SWITCH IT ON, OR HIT SMART AND LET THE STRIP BUILD ITSELF."*

The Fifteen Modules — In Depth

Every module editor gets the universal **OUT** knob (that stage's output trim, ±12 dB) and shows a **BYPASSED** badge when its stage is off. Ranges below are the real engine ranges; defaults in bold where they matter.

INPUT — the reference

What it does: Sets the level entering the strip, and carries the strip's input meters. Always first in the chain.

Controls

- **TRIM** (±24 dB) — input trim. Signed and smoothed; when AUTOGAIN/calibration is on, the ladder writes it for you.
- **POLARITY** — Ø flip. Rides the same ramp as the trim, so flipping is click-free.
- **SWAP L/R** — swaps the channels (20 ms crossfade).
- **PHASE — TIME** (0–50 ms) — the align delay. PHASE ALIGN writes it, and you can nudge it by hand. Delay only, so nothing ever plays early, and no latency is reported.
- **PHASE — ROTATE** (20 Hz – 2 kHz, **OFF** at the bottom of the travel) — a continuous phase turn, for the case polarity and delay together still can't fix. PHASE ALIGN solves it for you; see that chapter. Bit-transparent at OFF.

Both PHASE wells sit *outside* the INPUT card's on/off, so switching INPUT off can never slide you back out of phase.

On the card: live L/R input meters, "INV" / "L<>R" markers, "ROT" when a phase rotation is running, and the calibration micro-state.

Pro tip: Feed the strip something honest. The whole calibration story keys off the input peak — one LISTEN or CALIBRATE pass here and every threshold downstream starts telling the truth.

LOW-CUT — the rumble gate

What it does: A resonant high-pass that is a *true Butterworth at every slope* — -3 dB at the corner whether you choose 6, 12, 24 or 48 dB/oct, maximally flat passband, no surprises when you change the slope.

Controls

- **FREQ** (20–500 Hz, default 30) — the corner.
- **SLOPE** — 6 / **12** / 24 / 48 dB/oct.
- **RESONANCE** (0–1, default 0) — lifts a peak at the corner. 0 is exactly flat.

In your music: DC and rumble off everything; 80–120 Hz off vocals; tightening a boomy guitar. With RESONANCE up, a low-cut becomes a low-end *shape* — cut at 40 Hz with a bump at the corner is the classic club-kick trick.

Pro tip: The editor's curve is drawn from the same math the audio runs — what you see is measured to agree with what you hear to a tenth of a dB.

GATE — bleed, tails, and silence that breathes

What it does: A gate / downward expander with a transient-aware trigger. Cleans bleed and tails without chattering.

Controls

- **THRESH** (-80…0, default -28, calibrated-referenced) and **HYSTERESIS** (0–24 dB, default 6) — where it opens, and how much quieter the signal must get before it closes. Hysteresis is the anti-chatter control.
- **RANGE** (-80…0, default -80) — how far down "closed" is. -80 is silence; -12 is a polite expander.
- **RATIO** (1–50, default 10) and **KNEE** (0–18 dB, default 3) — how hard the expansion is below threshold.
- **ATTACK** (0.05–50 ms, default 0.5), **HOLD** (0–500 ms, default 10), **RELEASE** (5–2000 ms, default 80), **RELEASE MODE** (Linear / **Dual-Stage** / Auto).
- **ANALOG** (0–1, default 0.35) — analog-style character on the ballistics.
- **TRIGGER** (Level / Transient) with **ONSET** (0–1, default 0.35) — Transient mode fires on onsets rather than raw level, which is what keeps ghost notes alive on a drum gate.
- **DETECTOR** (Peak / RMS), **DETECT BAND** (Full / Low / Mid / High / Low+Mid) and **FOCUS** (±1) — band-limit and tilt what the gate listens to, so a kick gate doesn't hear the hats.
- **DRIVE** (0–24 dB) — drives the **detector only**: a sensitivity control that never touches your level.
- **SIDECHAIN** — Internal or External (the host sidechain bus), with SC HP/LP filters and **SC LISTEN** to audition the key.
- **LINK** (0–1, default 1), **MIX** (0–1, default 1), **LOOKAHEAD** (0–20 ms, default 0 — the one gate control that costs latency).

In your music: DI and amp bleed, tom tracks, noisy synth tails, tightening a live drum bus with DETECT BAND on lows.

Pro tip: For drums, TRIGGER = Transient with RANGE around -18 keeps the room under the hits instead of amputating it. Then HYSTERESIS is your chatter knob, not RELEASE.

EQ — six bands, all of them dynamic if you want

What it does: A six-band parametric where any band can become dynamic — moving only when its own detector crosses threshold — and any band can target Stereo, Mid, Side, Left or Right. On the mesh, a band's detector can even be *another track's* energy (that's UNMASK's carve — see the MIX chapter).

The curve: double-click empty space to add a band (the type is picked from where you click), drag a node (X = frequency, Y = gain), wheel for Q, double-click a node to remove it.

Per band

- **TYPE** — Bell, Low Shelf, High Shelf, High-Pass, Low-Pass (cuts ignore gain and dynamics). Defaults per band: Low Shelf / Bell ×4 / High Shelf at 80 / 240 / 700 / 1800 / 4500 / 11000 Hz.

- **FREQ** (20 Hz–20 kHz), **GAIN** (±18 dB), **Q** (0.1–8, default 1).
- **DYN** — makes the band dynamic: **THRESH** (-60…0, default -12, calibrated-referenced), **RANGE** (±24 dB — negative ducks, positive lifts), **ATTACK** (0.1–200 ms, default 10), **RELEASE** (5–2000 ms, default 120). The editor shows each dynamic band's live movement.
- **CHANNEL** — Stereo / Mid / Side / Left / Right.
- **LISTEN** — solo the band's own detector range.

Module-level: EQ OUT (±12 dB), **MIX** (0–1 — the whole EQ in parallel), auto-output option.

In your music: De-harsh at 3–6 kHz with a dynamic bell that only moves on the harsh notes. De-mud at 250 Hz on the Side only. A dynamic 100 Hz Side cut is a stereo tightener that doesn't touch your mono punch.

Pro tip: Static EQ for what's *always* wrong; DYN for what's *sometimes* wrong. If you find yourself riding a cut, that band wanted a threshold.

COMP — the honest compressor

What it does: A broadband compressor with three envelope characters and a threshold that — in its default mode — reads in **dB of gain reduction**.

Controls

- **THRESH** (-60…0, default -6) — see **THRESHOLD MODE** below.
- **RATIO** (1–20, default 2.5), **KNEE** (0–18 dB, default 6).
- **ATTACK** (0.1–200 ms, default 20), **RELEASE** (5–2000 ms, default 150).
- **ENV MODE** — **Clean** (textbook), **Opto** (a fast and a slow release stage blended — the smooth vocal glue), **Punch** (FET-style program-dependent release — the drama).
- **DETECTOR** — Peak / **RMS** / Hybrid. REDUCTION mode compensates for whichever you pick, using the crest the calibration actually measured on your material.
- **THRESHOLD MODE** — **REDUCTION** (default: -N = N dB of GR on the loudest material, any ratio) or **THRESHOLD** (classic, referenced to the calibrated level). Both read truest with the detector on Peak.
- **MAKEUP** (0–24 dB) + auto-makeup option, **MIX** (0–1 — parallel compression built in), **DRIVE** (0–24 dB into the compressor), **LINK** (0–1).
- **CHANNEL** — Stereo / Mid / Side (compress only the Mid, or only the Side).
- **SIDECHAIN** — Internal / External (host SC bus), detector-only SC high-pass (20–1000 Hz), **SC LISTEN**.
- **LOOKAHEAD** (0–5 ms, default 0).

The transfer curve in the editor is draggable — pull the knee point around and threshold/ratio follow.

In your music: 2:1 Opto on a vocal at -4 in REDUCTION mode is the whole "how much compression do I want" conversation in one knob. Punch on a drum bus. Mid-only compression to anchor a wide mix.

Pro tip: REDUCTION mode plus the GR readout on the card makes gain-staging conversations obsolete: decide the dB, set the dB, verify the dB.

DE-ESS — sibilance, surgically

What it does: Tames S's and HF harshness. In SPLIT mode it ducks only the high branch of a phase-exact band split — the body of the voice never moves. In WIDEBAND it rides the whole level, the classic broadcast way.

Controls

- **FREQ** (1–16 kHz, default 6 k) — where sibilance lives.
- **THRESH** (-60…0, default -18, calibrated-referenced), **RATIO** (1–20, default 4), **RANGE** (0–24 dB, default 12 — the most reduction it may ever apply).
- **ATTACK** (0.1–20 ms, default 0.5), **RELEASE** (10–500 ms, default 60).
- **MODE** — **Split** / Wideband. **CHANNEL** — Stereo / Mid / Side.
- **LISTEN** — auditions **what's being removed**. If you hear vowels in there, back off.
- **LINK** (0–1).

In your music: Vocals, obviously — but also cymbal wash on overheads (Side mode) and fizzy hi-hats in a top loop.

Pro tip: Set RANGE before THRESH. A de-esser capped at 6 dB can be set aggressively and still never lisp.

DUCK — the pump, three ways

What it does: Rhythmic or triggered level dips — sidechain pumping without a sidechain. One drawn dip shape, three ways to fire it.

Controls

- **TRIGGER** — **SYNC** (the dip rides the host's beat grid, phase-locked to the timeline: jump anywhere in the song and it's already right), **SIDECHAIN** (a hit on the host SC bus fires one traversal), **FOLLOW KICK** (the mesh: the dip keys off a KICK-typed instance's low band — see the MIX chapter).
- **AMOUNT** (0–100 %, default 100) — how deep. A percentage of the signal, not dB: **100 % is actual silence**.
- **LENGTH** — the note value of one cycle: 1/1 … 1/32 including dotted and triplet, default **1/4**.
- **CURVE** (±100 %) — the shape of the climb back up: scooped through linear to snappy.
- **CURVE LEN** (1–100 %, default 25) — how much of the cycle the climb takes. 25 % = open again a quarter of the way in.
- **TAIL** (default on) — dive back down at the end of the cycle instead of holding open, so consecutive dips land like a pump, not a gate.
- **SC THRESH** (-60…0, default -30 — absolute, it keys a foreign signal) and **SC LISTEN** — Sidechain mode's trigger level, and a way to hear the key while you set it.
- **BELOW** (100 Hz–20 kHz, default 20 k = full band) — duck only *below* this frequency. Leave the top of the track standing while the lows make room.
- **LOOKAHEAD** (0–20 ms), **MIX** (0–1), **OUT** (±12 dB).

The editor draws the actual dip — the same math the audio runs — with a beat-true playhead.

In your music: The house pump (SYNC, 1/4). A bass that steps out of the kick's way (FOLLOW KICK with BELOW at ~150 Hz). A pad that breathes under an unquantized live kick (SIDECHAIN).

Pro tip: BELOW is the difference between "pumping" and "mixing". Full-band ducking is an effect; ducking below 150 Hz is a low-end arrangement.

TRANSIENTS — attack and body, decoupled

What it does: A transient designer with *real windows*: it shapes the attack for a length you set, and the sustain region for another, with detection you can focus by band.

Controls

- **ATTACK** (±12 dB) and **SUSTAIN** (±12 dB) — the two headline moves: more or less hit, more or less body.
- **SENSITIVITY** (0–1, default 0.5) — how readily hits are detected.
- **ATTACK LEN** (5–200 ms, default 30) and **SUSTAIN LEN** (50–3000 ms, default 3000 = unlimited) — how long each region lasts.
- **CONTRAST** (0–1) — dynamic depth: strong onsets get shaped harder than weak ones, which keeps ghost notes human.
- **FOCUS** (±1) — tilt detection toward lows or highs.
- **ATK DRIVE / SUS DRIVE** (0–1) — saturate the attack or sustain stream on its own.
- **OUTPUT** (±24 dB), **MIX** (0–1), **MONITOR** — Mix / Transient / Sustain (solo either stream while you shape it).

In your music: Snap on a lazy snare; room out of an over-ambient kit (SUSTAIN down); pick attack out of a strummed acoustic without EQ.

Pro tip: MONITOR = Transient while you set ATTACK LEN — you'll *hear* where the hit ends and the body begins, then set the window right at the seam.

MB-COMP — three bands of glue

What it does: A three-band multiband compressor — the full per-band control set, not a macro.

Controls

- **CROSSOVERS** — LOW/MID at 50–1000 Hz (default 200), MID/HIGH at 500–8000 Hz (default 2000). The split reconstructs flat when the bands do nothing.
- **Per band (LOW / MID / HIGH)** — THRESH (-48…0, default -12, calibrated-referenced), RATIO (1–12, default 2), KNEE (0–18, default 6), ATTACK (0.1–200, default 10), RELEASE (5–2000, default 120), MAKEUP (0–24), DETECTOR (Peak / RMS / Hybrid), ENV (Clean / Opto / Punch), per-band **BYPASS** and **SOLO**.
- **MIX** (0–1), **DRIVE** (0–24 dB in), **OUTPUT** (±24 dB) + auto-makeup option.

In your music: Low-band-only control on a bass that moves too much; de-boxing a drum bus by compressing just the mids; low-band glue on a 2-bus. SMART SETUP never sets this module up — no profile does — so what you dial in here is yours and survives every learn.

Pro tip: Solo a band, find the problem, un-solo, compress. Multiband set "because it's there" is how mixes go grey — one working band usually beats three busy ones.

STEREO — width with a safety net

What it does: An M/S stereo imager with bass management, optional three-band width, and a governor that stops width from becoming phase trouble.

Controls

- **WIDTH** (0–2, default 1) — 0 folds to mono, 2 doubles the side.
- **BASS MONO** (0–500 Hz, default 0/off) — mono below the corner. The club-safe switch.
- **3-BAND** + **LOW / MID / HIGH widths** (each 0–2) with crossovers (default 250 / 2500 Hz) — widen the top without touching the low mids.
- **MONOIZE** (0–1), **MID GAIN / SIDE GAIN** (±12 dB), **BALANCE** (±1).
- **GOVERNOR** — ON / THRESH (±1) / AMOUNT (0–1, default 0.5): a dynamic width limiter that reins the image in when correlation goes bad, and leaves it alone otherwise.
- **MIX** (0–1), **OUT** (±12 dB).

Metering: live correlation, width %, balance and mono-fold energy on the card.

In your music: Widening pads and overheads; BASS MONO at 100–150 Hz on any wide bass; SIDE GAIN as a subtler wide/narrow than WIDTH itself.

Pro tip: Width you notice in solo is usually too much in the mix. Set it in context, watch the correlation dot on the OUTPUT page, and let the GOVERNOR carry the worst case.

REVERB — space that steps aside

What it does: An algorithmic reverb with per-band decay — and a built-in **wet ducker**, so the tail makes room while the dry signal plays and blooms in the gaps. Space without wash.

Controls

- **MIX** (0–1, default 0.1) — the reverb runs wet-only internally; this is the blend.
- **DECAY (T60)** (0.1–20 s, default 1.4) and **SIZE** (0.1–4, default 1).
- **DAMPING** (0–1, default 0.5) and **DIFFUSION** (0–1, default 0.707).
- **MOD DEPTH / MOD RATE** (0–1 / 0–5 Hz) — tail modulation, from still to chorused.
- **LOW / MID / HIGH DECAY** (0.1–4× each) with **LOW CUT / HIGH CUT** corners (default 500 Hz / 4 kHz) — a short low tail under a long mid tail is how you get "big but clean".
- **DUCKING** (0–1, default 0) — the wet ducker: keyed by the dry signal, 8 ms in, 250 ms back. The editor shows its live gain reduction.
- **OUT** (±12 dB).

In your music: Vocal space that never blurs the words (DUCKING at 0.3–0.5). Snare rooms. SMART SETUP seeds tasteful, ducked settings on the profiles that want space.

Pro tip: Reach for DUCKING before you shorten the decay. A long, ducked tail sounds bigger *and* clearer than a short, constant one.

DELAY — echoes on the grid

What it does: A tempo-synced stereo delay with a filtered, saturating feedback loop, ping-pong, cross-feed, tape-style wobble — and the same wet ducker as the reverb.

Controls

- **SYNC** (default **on**) + **DIVISION** (1/1 … 1/32, dotted and triplet, default 1/8) — or **TIME** (1–2000 ms, default 350) free-running. Synced time follows tempo changes live, ramps included.
- **FEEDBACK** (0–0.95, default 0.25) — repeats. It cannot run away: stability is a tested guarantee.
- **MIX** (0–1, default 0.1).
- **HP / LP** (20–2000 Hz, default 100 / 1–20 kHz, default 8 k) — filters *in the loop*: every repeat gets darker and thinner, like a real machine.
- **PING-PONG**, **WIDTH** (0–2), **CROSSFEED** (0–1), **DRIVE** (0–1 — feedback saturation).
- **MOD DEPTH / MOD RATE** (0–1 / 0–5 Hz, default 0.3) — wow and flutter on the line.
- **DUCKING** (0–1) — the echoes step back while the dry plays, bloom in the gaps.

In your music: The 1/8-dotted vocal throw that never treads on a word (DUCKING at 0.5–0.7). Ping-pong on a pluck. DRIVE plus tight HP/LP for the dub tape loop.

Pro tip: SMART SETUP always writes delays synced, with the division picked per profile — RAP gets 1/8, synths 1/8-dotted. Change the division, not the sync.

MOD — the movement rack

What it does: Five classic modulations over one LFO: **CHORUS, FLANGER, PHASER, TREMOLO, AUTOPAN** — free-running or locked to the host grid.

Controls

- **MODE** — the five above (default Chorus). They share only rate and depth; everything else re-homes per mode.
- **SHAPE** — Sine / Triangle / Square / Saw / Random step.
- **SYNC** (default off) + **DIVISION** (1/32 … 1/1, dotted/triplet) — grid-locked LFO phase, derived from the timeline: two synced instances move identically. **RATE** (0.05–10 Hz, default 0.8) free-running.
- **DEPTH** (0–1, default 0.5) — delay swing, notch spread, or gain swing depending on mode.
- **MIX** (0–1, default 0.5) — for the wet-path modes. **Tremolo and auto-pan ignore it** — they modulate level, not a wet path, which is also why they're attenuation-only and safe anywhere in the chain.
- **SPREAD** (0–1, default 0.5) — LFO phase offset L/R, 0–180°: the width of the movement.
- **FEEDBACK** (±0.95) — flanger/phaser resonance; negative inverts the comb.
- **NOTCHES** (2–12, default 4) — phaser thickness.
- **OUT** (±12 dB).

The editor draws one LFO cycle from the live settings, with a playhead on the engine's real phase.

In your music: Chorus on backing vocals; a synced square tremolo as a gate-substitute; auto-pan on percussion beds; a slow saw phaser for the sound-design pass.

Pro tip: SYNC a triangle AUTOPAN at 1/2 and it's an arrangement move, not an effect — the part changes sides every bar, locked to the song forever.

WARMTH — the analog stage

What it does: Six saturation models with bias, tone, and — on the Tape model — hysteresis and wow. Independently oversampled, so it can sit anywhere in the chain without paying for the whole strip.

Controls

- **DRIVE** (0–1, default 0 — the stage ships **on** but clean) and **MIX** (0–1, default 1).
- **MODEL** (1–6) — from tube through tape, transformer-style grit to console glue. Internal auto-gain holds the level as you drive, so DRIVE changes the *sound*, not the volume.
- **BIAS** (0–1) — asymmetry: the even-harmonic warmth.
- **TONE** (±1) — tilt the saturated signal darker or brighter.
- **TAPE HYST** (0–1, default 0.5) and **TAPE WOW** (0–1) — the Tape model's memory and wobble.
- **MULTIBAND** — split the saturation per band.
- **OS** (1/2/4/8×, default 2×) — the stage's own oversampling.
- **OUT** (±24 dB).

In your music: Console glue at DRIVE 0.1 on everything; tape on a synth bus with a breath of WOW; transformer grit on drums.

Pro tip: If you want *cleaner* loud saturation, raise this stage's OS — and see OVERSAMPLING QUALITY in the Specs chapter for the Fast/Clean switch that governs how the strip decimates.

LIMITER — the last word

What it does: A lookahead brickwall limiter, pinned last — with a real **anti-aliased clipper built in**, so one knob decides how much of the peak work gets clipped off before the wall. The ceiling is a guarantee: nothing leaves the strip above it. Switch **TP** on and that guarantee extends to inter-sample peaks.

Controls

- **CEILING** (-24...0 dB, default -1) — absolute, always.
- **GAIN** (0–24 dB) — drive into the limiter.
- **CLIP** (0–100 %, default 0) — the built-in clipper: how much of the peak work the clipper takes before the limiter. At 0 % it's parked out of reach — bit-inert. Turn it up and a clip ceiling slides down from 12 dB above the limiter ceiling toward the ceiling itself; at 100 % the clipper shaves every over and the limiter just mops up. It runs anti-aliased in the limiter's own oversampled domain, tracks CEILING and GAIN, and adds **zero latency**. Clipped transients read as punch; limited ones read as control — this knob is the blend.
- **CHARACTER** (0–1) + **TYPE** (TRANSFORMER / TUBE / TAPE) — colour on the way in.
- **LOOKAHEAD** (0–20 ms, default 1.5), **ATTACK** (0.01–10 ms, default 0.05), **RELEASE** (1–1000 ms, default 80), **RELEASE MODE** (CLEAN / DUAL / **AUTO**).

- **LINK** (0–1, default 1), **CLIP SAFETY** (OFF / **SOFT** / HARD) — a separate base-rate safety net after the oversampling, guarding the ceiling against reconstruction ripple. Not the CLIP knob: SAFETY is the seatbelt, CLIP is the sound.
- **OS** (1/2/4/8×, default 1×).
- **TP** (default off) — true-peak limiting: inter-sample peak detection at 1/8-sample resolution drives a real gain duck, so the ceiling holds as a **true-peak** ceiling at any oversampling setting. Costs ~2 ms of latency while on.
- **S** (delta-solo) — the limiter now has one: hear the clip, the push and the wall together — exactly what the last stage is costing you.

The editor shows gain-reduction history and the true-peak meter against the ceiling; while the clipper is working, a warm **"CLIP -X.X DB"** readout appears on the trace. The GR number is the total peak work — clipper, limiter core and TP stage together.

In your music: -1 dBTP with TP on for anything going to streaming. CLIP at 30–60 % on drums and full mixes for loudness that stays punchy. CHARACTER up for the "loud" record; CLEAN/AUTO release for the transparent one.

Pro tip: Push GAIN, then ride CLIP while delta-soloing the limiter (S). When the delta is transient ticks, the clipper is shaving cleanly; when it grows body, the wall is working — pull CLIP back or drop GAIN. With TP off the ceiling is sample-accurate — if the number on the meter is a delivery spec, switch TP on and forget about it.

Output & Gain Staging

The XY pad

Always on screen at the right. **Vertical = output gain (±12 dB), horizontal = pan**; double-click resets. It's a *relative-pickup* pad — grabbing it never jumps the value — and SHIFT is fine control. Two padlocks:

- **LOCK CENTRE** — pan pinned to centre; the pad is gain-only.
- **LOCK VOL** — gain pinned; the pad is pan-only.

The pad is truly the **last** gain in the strip: the calibration lands the chain at the reference, and the puck is a clean offset on top. Under the puck, the final output's stereo field breathes as a particle cloud — ambience, not a meter.

The OUTPUT big view

Click the OUTPUT rail (or TAB past the last module). It answers two questions:

"How loud am I, and am I safe?"

- **OUTPUT SPECTRUM** — strip input vs finished output.
- **Loudness column** — **MOM / SHORT / INT** LUFS, **RANGE** (LRA), **TRUE PK** (dBTP, held), **CORR** with a phase dot. All read post-everything. A silent channel prints "--", never a fake zero. **Click the column to reset** the held figures — integrated restarts, the true-peak hold clears.
- **L/R meter bridge** with a **-14 LUFS** reference horizon.

"What did the strip do to my level?"

- **The gain map** — one line, in signal order: **IN** (absolute) → each enabled stage's own ± contribution ("EQ +2.4") → **TRIM** (the strip's own correction) → **OUT** (absolute). *Ends are levels, middle is change.* Figures latch slowly so you can actually read them.
- **GAIN CORRECTION** — the final trim the calibration set (draggable to overrule, double-click for none; shows "MEASURING..." while a window runs).

The action rail

- **CALIBRATE** — the AUTOGAIN switch (same states as the header chip).
- **RE-MEASURE** — run the ladder again on the current audio.
- **DIM** — temporary monitor cut.
- **SET** — the GAIN-STAGING SETTINGS popover: **Target** (-24...-3, default -12 dBTP), **Window** (0.5–6 s, default 2.5), **Range** (3–24 dB, default 24), **RESTORE INPUT VOL** (calibrate back to the *input's* level instead of -12), and **INIT THE_STRIP 3** — a two-click-armed full reset ("! CLEAR EVERYTHING?"), one undo step.
- **BYPASS** — the whole strip off. The strip prints its input untouched, reports zero latency so the host re-aligns, and the meters keep reading so the window never looks dead.
-

MATCH (default **on**) — level-matched bypass: the strip measures its own level change and plays the dry signal at the processed level (up to ±24 dB of correction), so an A/B compares **tone, not loudness**. Switch MATCH off for a straight, bit-exact null bypass. The first bypass explains itself once: *"Bypass is level matched — you hear the dry signal at the processed level."*

The toolbar power button is the third, bluntest layer: a hard bypass at unity with a full-screen overlay — "BYPASSED - CLICK ANYWHERE TO RE-ENABLE PROCESSING".

The MIX View — one mixer for the whole session

Press **M** or flip the header toggle to **MIX**. Every THE_STRIP 3 and MINI instance in the session appears as a channel card — automatically. No routing, no setup: the instances find each other the moment they exist.

The channel card

Top to bottom: **name** (double-click your own card to rename it — the name broadcasts to every window; the local instance wears a cyan dot; MINIs wear a "MINI" badge), **signal-type chip**, a **mini-rack pip row** (one pip per stage, lit in the stage's colour when enabled — the whole strip's state at a glance), the **masking status line + UNMASK pill + heat ribbon**, a **stereo meter**, the **fader** (with a standing dB readout), **pan**, **M / S** chips, the **DUCK** chip, **EDIT STRIP**, and the **LUFS / TP / COR** readout block.

Gestures: relative-pickup fader with SHIFT fine and a floating value tag ("-6.0 dB", "L 35"); double-click fader = unity, pan = centre; right-click the fader to type a value ("-6", "-inf"); **click the readout block to reset** that instance's integrated LUFS and true-peak hold — works on remote instances too; wheel scrolls the mixer; drag a name bar to reorder the cards.

The type chip is the card's identity and its learn button: click to pick the signal type (the picker carries the same search as SMART SETUP — type to find a profile, or a word that fits one), **alt-click (or right-click) to run a LISTEN on just that card** — headlessly, without opening its window. During a learn it becomes a progress bar ("LISTENING 37%", amber "WAITING - PLAY" when starved); a learn that gave up latches an amber **NOT SET UP** until the next one. The picker's footer carries the **AUTO-MIX lock** ("LOCK FROM AUTO-MIX" / "LOCKED - AUTO-MIX KEEPS ITS HANDS OFF") — see the AUTO-MIX chapter; a locked card wears a small warm padlock on its type chip.

Solo and mute are session-wide, honest, and routing-aware **in both directions**: solo mutes the others, remembers every mute you had set by hand, and restores them when the solo clears. A bus that carries the soloed track stays up — and **soloing a bus keeps everything that feeds it unmuted** too, so a soloed drum bus actually receives its drums. A solo-imposed mute paints the M chip in solo gold — visually distinct from your own warm-red mutes — and the status line says **MUTED BY SOLO**. When the mesh can't tell what feeds what yet, it errs open — a solo that is momentarily too generous beats one that plays silence. **Alt- or Cmd/Ctrl-click any S chip clears every solo** ("Solo cleared — every channel is back to the mute you set."). Solos never survive closing the editor.

Groups — the mixer organises itself

Tracks that belong together become a group — mostly by themselves. The moment two or more tracks share a family, they group automatically: **VOICE, DRUMS, PERC, BASS, GUITAR, KEYS, SYNTH, ORCH, FX**. Grouping alone never changes your audio — it's the group *gestures* that do things.

- **The group chip** sits at the right end of every track card's name row — a filled letter in the group's colour ("D" for DRUMS), or a quiet dashed square when ungrouped. Click it to choose the track's group, start a **NEW GROUP** ("START ONE WITH THIS TRACK"), or pick **NO GROUP** ("AND DON'T AUTO-GROUP IT AGAIN") — a remembered opt-out the auto-grouper respects forever. Buses and the master never group.
-

The group band floats above the run of member cards: name and count ("DRUMS - 4", "- LINKED" when linked), a **LINK** chip, **M** and **S** chips, and a fold glyph. Group **M** mutes every member but *remembers which mutes were its own* — un-pressing releases only what the group muted, never a mute you set by hand. The M chip only half-lights when just some members are muted — it never claims more than it did. Group **S** solos the lot (with the same routing-aware semantics as any solo). Right-click the band for the group menu: COLLAPSE / LINK / **RENAME** / **UNGROUP ALL** ("EVERY TRACK BACK ON ITS OWN").

- **Collapse** folds the run into one card: group name, count, a stereo meter with an approximate summed loudness ("~-18.2" — the ~ is honest), M / S, LINK, and a **shared fader**: drag it and every member moves by the same dB — a VCA ride with your offsets preserved (a silent member is never resurrected, and if a track hits the rail the mixer says so once: *"A TRACK HIT THE FADER RAIL - KEEP GOING AND THE GROUP'S OFFSETS WILL CHANGE"*). Click the card to expand; drag it to carry the whole group as one block.
- **LINK** is the ganged-settings switch: *"LINKED - LOW CUT, COMP TIMES AND DUCK LENGTH FOLLOW WHOEVER MOVES THEM. VALUES MATCH ON FIRST TOUCH, NOTHING CHANGES UNTIL THEN."* Exactly five parameters gang — the LOW-CUT frequency, the compressor's ATTACK and RELEASE, and the duck's LENGTH and CURVE LEN — and they follow the touch: move one of those knobs on *any* member's strip and the same value lands on every member. Levels, EQ and thresholds never gang. While a strip is in a linked group, its ganged knobs wear a small **LINK** pill so you always know which moves are group-wide.
- **RENAME** gives a group your own name — up to 28 characters, edited in place on the band. The field opens on *your* name and empty when there isn't one yet, with the automatic name showing behind the caret as a hint: pressing RETURN on a field you never touched would otherwise pin "DRUMS" as a name you chose, and it would stop following the family from then on. Clear the field and the automatic name comes back. The name reaches every instance in the session and comes back with the song; dissolve a group and its name goes with it, so the next group to take that slot is born nameless. Renaming is only a label — unlike LINK it doesn't fix who is in the group.

Groups, their names, collapse state and the opt-out all persist with the session.

BYPASS ALL (header) is the whole-session A/B: one click and every strip in the mesh comes out — the processing, and the mixer's own faders, pans and trims with it, because AUTO-MIX's entire balance lands on those too. Your mutes and solos are held, so you keep comparing the same set of tracks, and one shared offset — worked out once at the click and then frozen — holds the sum at its level, so what decides the comparison is the mix, not the loudness. The button burns warm red — the only red on the screen — while the mesh is bypassed. One limit worth knowing: a track panned harder than about 94 % cannot be brought all the way back to centre, so it sits a little off-centre while bypassed.

Buses and masters know what they are

The mesh works out your routing by listening: an instance whose signal *contains* the others is recognised as a **BUS** or **MASTER**, pinned to the right of a hairline rule — everything left of the rule adds up to everything right of it. Role cards show "MASTER" or "BUS - SUM OF n"; they carry no M/S chips and sit outside the UNMASK order (every track is already inside them). The bus card's well tells **that bus's own story**: **"3 MAKING ROOM BELOW"** when three of the tracks it carries are currently holding carves, **"NO CARVES BELOW"** when none are. You can override the detector: click the role block to

hold a card as a TRACK, click "HELD AS TRACK" to hand it back. Explicit picks and bus/master LISTEN profiles always outrank the detector.

Every bus type is really a bus. Declaring a card VOCAL BUS, DRUM BUS, GUITAR BUS, SYNTH BUS, ORCHESTRAL BUS, PERCUSSION BUS, KEYS BUS, BASS BUS, FX BUS or GROUP BUS makes it behave as one everywhere — it group-solos, sits out of the pecking order, is never carved, and learns in the bus tier of an AUTO-MIX run. (A whole kit recorded on one *track* is the **DRUM KIT** profile instead — it ranks and carves like a track.) And a bus on AUTO **inherits the family of what it carries**: a generic bus measurably summing your drum tracks quietly becomes a DRUM BUS — majority-female vocal children even land it on VOCAL BUS (FEMALE). Your explicit picks are never overridden.

EDIT STRIP — drill into any channel

EDIT STRIP opens that channel's *entire strip* below the mixer — the full rack, every editor, every trim, live visualisations fed by the **remote instance's own audio**. Edit anything; changes land on the remote instance instantly. Calibration commands route to the drilled channel. Your local undo history is parked at the boundary and handed back when you leave. Two honest limits: LISTEN always captures the *local* instance's audio, so the strip refuses to learn a remote channel from here — use the type chip's alt-click learn, or that channel's own window — and drilling into a MINI shows the five stages a MINI actually has.

HYBRID — the mixer as a rail

The third view: the mixer collapses to a one-row rail above the open channel's full strip. Rail cards carry name + dB, a horizontal fader pill with a hairline meter, and M/S. Click a card to switch which channel is open below; click the open card (or ESC) to close the drill. The bright card is the channel you're editing — everything else recedes.

UNMASK — decide who wins

Masking is the mix problem: two tracks fighting for the same frequencies at the same time. UNMASK turns it into one decision you can drag.

The pecking order

Every track in the mesh sits in **one total order**. Rank #1 is the most important track in the song. A track only ever steps back where something **above** it collides with it — and only **while that something plays**, because the carves are dynamic EQ bands keyed by the winner's actual energy in that band, live across plugin instances. No sidechain cables; the mesh *is* the sidechain.

The order seeds itself from signal types the moment strips are typed — lead vocals up top, beds below — so the feature works before you touch it. From then on the order is yours.

- **The card pill**: the top track wears **#1**; every other track wears **UNMASK**. Pressing it means *"I want to hear THIS"* — the track goes to the top and the mesh re-derives who steps back for whom. The status lines across the mixer tell the story in one vocabulary: **"STEPS BACK FOR LEAD VOX"**, **"2 STEP BACK"** (on the winner), **"UNDER GUITARS"**, **"CLASH VS SYNTH"**, **"CLEAR"**.
- **The heat ribbon** on each card maps *where* the collisions are — warm where it's contested, with a hairline at the dominant collision frequency.
- Collision detection is **pan-aware** (hard-opposite tracks don't collide), perceptually weighted (the 1–4 kHz region counts most; sub-30 Hz doesn't), and averaged over ~1.6 s so it reads the arrangement, not the moment. Home-turf bonuses mean bass wins the lows and vocals win the presence region even between near-equal ranks.
- **The carves are precision-fit.** Each carve's width is *measured*, not fixed: a broad 200–500 Hz mud smear gets one wide, gentle bell; a narrow honk gets a surgical notch. Depth is sized by how much is actually masked — low-mid mud finally gets cuts as deep as it deserves — and two genuinely separate collisions (a 300 Hz smear *and* a 700 Hz honk) can both be answered on the same track.
- **Kick and bass are never EQ-carved** — not against each other, and not for anything above them either: the foundation is what the rest of the mix makes room *for*. The kick/bass pocket belongs to the DUCK: promoting a kick arms follow-kick ducks on the bass-family tracks instead. Ask a kick to duck itself and it refuses, in words.

The UNMASK board

The header **UNMASK** chip (its badge counts the tracks currently carving) opens the board: **"DRAG TO SET WHO WINS — EVERYTHING BELOW STEPS BACK FOR IT."** Rows are the roster in order; drag to re-rank. Big sessions scroll — the wheel, or the slim gutter beside the rows, and a quiet "+N MORE - SCROLL" tells you when there's more below the fold. Changing the order *is* the action — there is no apply button. Each row reports its state in plain language: **"-4.2 dB · WHEN LEAD VOX PLAYS"** on a bending track, **"3 TRACKS STEP BACK FOR IT"** on a winner, **"WAITING FOR LEAD VOX"** when armed but not yet keyed, **"DUCKED BY KICK"**, **"MASTER - EVERY TRACK IS INSIDE IT"** below the rule.

- **The lock** (per row) — never carve this track, wherever it ranks.
- **HOW HARD** — one dial scaling every carve depth: GENTLE / MEDIUM / STRONG.

- **UNMASK on/off** — releases the whole layer *without forgetting the order*.
- **COMPARE** — hear the mix without UNMASK: hold it for as long as you want to listen, or give it a quick click and it stays off until you click it again. It suspends only the carve bands, mesh-wide, crossfading them out so nothing clicks, and the mixer says so while it holds — **UNMASK BYPASSED** on the cards, the carving rows dimmed. It never touches the strips, their levels, your automation or the saved session — a true isolated A/B of the feature itself.

Order changes are one cross-instance undo step. Carves land on each strip's free dynamic-EQ bands — SMART SETUP always leaves **at least three** of the six free for them (it spends its budget on the strongest three tonal moves instead), so a busy track can hold three simultaneous carves, up to four when fewer bands are engaged, and UNMASK never repurposes a band you set by hand. Open a carved strip's EQ and you can *see* the network's bands working. Pull a masker's fader down or mute it and its carves release: the mesh follows what's audible, not what's playing. Everything — ranks, locks, amount, engage — persists with the session and survives reopening it.

THE MIX MAP — your session as a room

A column of faders tells you what you *set*. It never tells you where anything actually is, and a pan knob per card means holding a dozen knob angles in your head. **MAP** — in the mixer header, on the same run as UNDO and UNMASK — opens the whole session as a stage seen from the audience.

STAGE

Across is the pan. The higher a track's fader, the nearer and the bigger it stands — **+6 dB at the front edge, -24 dB at the back wall**. There is no legend to learn: the loudest thing in your mix is the big one at the front, and the track you forgot about is the small one in the corner. **CLASH** marks the tracks standing on the same spot as each other.

CAMERA picks the view — **FRONT**, **RAISED** or **SIDE** — and **DRIFT** sways it slowly for when you just want to look. Drag the empty room to turn the stage yourself; double-click empty space to bring the camera back.

The board opens READ ONLY, on purpose. A fader you nudge here doesn't just move a fader — it lands on the mix ledger as *your* offset, which means every AUTO-MIX run from then on measures around it and faithfully puts it back. One slip would become a permanent opinion about the session. So the board opens safe: the amber **EDITING** chip arms it, and it disarms again the moment you close it.

Armed, one drag moves both things at once. Across is the pan, towards you is louder. The first bit of travel decides which of the two you're doing and then holds it — on a stage both values come off one pointer position, and a free drag would quietly repan the vocal you only wanted louder. **ALT** frees both axes, **SHIFT** is fine, and **double-click** a track sets its fader to 0 dB. Every move is one mesh-wide undo step, and a click that changes nothing leaves no undo entry at all.

What isn't in the room is listed under it. The **NOT HEARD** tray names every muted, solo-muted, bypassed and silent track, grouped by the reason — so "which track am I not hearing, and why" is a question the board answers instead of one you go hunting for. Buses and the master don't stand among the tracks, because a bus already contains them; they're named under the room with their level, as what it all **ADDS UP TO**.

FREQUENCY

The second view on the same board, answering a different question: who owns which part of the spectrum. Low to high across, the outline is the whole mix, and each layer is one track's share of it, stacked by where that track's energy actually sits. Colours from the same family alternate light and dark down the stack, so a drum session reads as separate tracks instead of one red mass. Under the well, **LOWS / MIDS / HIGHS** name who owns each third.

AS HEARD draws the mix with UNMASK's carves in it. **BEFORE UNMASK** draws it without them, the other side dashed in behind — so you can see exactly what the network took out to make room for the vocal. That's a view toggle, not audio: nothing switches off while you're looking.

Empty states tell you what to do rather than what's missing — *"NO SPECTRUM YET - PLAY THE SESSION AND EACH TRACK FILLS IN ITS SHARE."*

Pro tip: open the map right after an AUTO-MIX run, before you touch anything. A whole run's balance is one look, and the tracks the board flags as standing on the same spot are usually the ones worth a REVISIONS chip.

AUTO-MIX — the one-button pass

The violet header button reads what it would do before you press it:

- **AUTO-MIX** — everything needs setting up.
- **AUTO-MIX +N** — incremental: N strips are new (or re-typed) and would be set up; everything already set up keeps its settings and is only re-levelled where needed.
- **RE-BALANCE** — nothing to set up; a level-only pass over every strip and the mix trim.

Alt-click is the from-scratch run — "I've made a mess, decide the whole thing again": every strip re-learned, the pecking order re-seeded from types — unless you have authored the order yourself by dragging a row on the UNMASK board (over typed tracks): an authored order, like an authored fader or pan, survives every run.

Preflight

A press that would change anything first asks on a small card:

- **MAKE IT WET?** — ducked reverb + delay for the sources that want space (vocals, keys, snares…). **Kick, bass and spoken word stay dry** — the wet decision is type-aware, per strip.
- **FOUR TO THE FLOOR?** — kick on every beat? Then the ducks pump on the host grid at 1/4 instead of following the kick.
- **SPREAD THE STAGE?** — pans supporting tracks out symmetrically. **Kick, bass and leads stay in the centre** (see below).
- **HOW LONG SHOULD IT LISTEN?** — 3 / 6 / 12 s per strip. Longer hears more of the arrangement — and sets levels off a longer window too; every learn and level window in the run scales with it.

RUN AUTO-MIX starts it; clicking outside is "not now". Your answers become the defaults for next time (this session).

One thing to do before you press it: set every DAW track volume to 0

AUTO-MIX needs every DAW track volume at 0. This is the one requirement it has, and it is worth understanding rather than just obeying, because nothing in the plugin can work around it.

A strip sits in the insert chain, **upstream of your DAW's track volume**. The only thing a host ever tells a plugin about its track is the **name** and the **colour** — there is no volume in that message, in any format, in any DAW. So every level the mesh knows is the level *before* your track volume, and every sum it reasons about assumes those volumes are at 0. Leave one at −8 and AUTO-MIX is balancing a mix that is not the mix you are hearing: the relative decisions (SECTIONS, the stem placement, the 2-TRACK ratio) are all made on the wrong proportions.

Do your gain staging **inside the strip** instead — the MIX view's own faders are what AUTO-MIX reads, writes and remembers, and the mix ledger is built on them. That is what they are for.

When it can tell, it tells you. With a strip on your master, the run compares what that instance actually receives against what your loudest single track is putting out. A sum cannot be quieter than its loudest member, so if less is arriving than one track sends, something between them is pulling level away —

and the report says so, with the size, and asks you to set the track volumes to 0 and run it again. It is a deliberately conservative check: it will never accuse a session whose track volumes *are* at 0, and the price of that certainty is that it only catches the gross case. A couple of dB of drift will pass unnoticed, so the requirement is still yours to keep.

The stage spread (and the human vocal rule)

Say yes to SPREAD THE STAGE? and the run also *places* the mix: the foundation — kick, bass, 808, sub, snare, a DRUM KIT, every bus and **the** lead vocal — stays dead centre, and the supporting cast fans out symmetrically from the middle by pecking rank (12 / 20 / 28 / 36 %), alternating sides per family so the stage doesn't lean. Vocal doubles and backings go wider (40 / 50 %). Nothing the run places on its own goes past halfway out — a one-button mix that throws parts at the walls is a mix you undo, and going wider is one drag, or one GUITARS WIDER click, away. The **human vocal rule**: only the best-ranked lead voice keeps the centre — a session with two "lead" vocals doesn't get two centred leads, it gets a lead and a wide double, because two centred leads isn't a decision, it's an unresolved fight. A track whose *name* picks a side ("GTR L", "OH RIGHT") keeps that side — the spread only chooses the width. Incremental runs never reshuffle a stage you've arranged; one Undo takes the whole placement back with everything else.

Two tracks — the vocal-to-bed ratio

A vocal over an instrumental is the one session a per-track balance cannot get right on its own: every strip lands on the same −12 dB reference, so the loudness gap between a compressed voice and a limited beat is nothing but their crest-factor difference — never a decision anyone made. So when a run finds **exactly one voice and one bed** (a vocal — or a VOCAL BUS — next to an INSTRUMENTAL, an AUTO track, a DRUM KIT, or a single instrument like a guitar or piano; a master above them is fine, a kick or a bass next to the voice is not a duet), it adds an act after that tier has settled: **2-TRACK**. It listens to both cards for ~6 s (longer with a longer listen length), measures their real loudness — gated so the vocal's gaps don't count as quiet vocal — and moves **only the bed's fader** until the vocal sits where the two profiles say it should:

- the base is **a LEAD VOCAL 4 LU under a POP bed** (the full mix then comes out about 1.5 LU louder than the bed alone);
- the vocal's profile sets its forwardness — **RAP VOCAL** 1.5 LU closer, **SPOKEN / PODCAST** 10 LU *over* its music bed, backing vocals, doubles and choirs 4 LU further back;
- the bed's genre nudges it — **HIP-HOP / TRAP** and **ACOUSTIC** beds carry the voice 1 LU further out, **BAND / ROCK** 2 LU and **ELECTRONIC / DANCE** 3 LU deeper in.

The voice always leads the pecking order (the bed is ranked under any voice by construction), so the bed also holds the UNMASK carves — it steps back only where and while the voice sings. A move of 3 dB or more re-plans those carves against the new balance before the run goes on. The ledger says what happened in numbers — *"I pulled BEAT down 3.0 dB so VOX sits 3.0 LU under BEAT — the LEAD VOCAL over an INSTRUMENTAL (HIP-HOP / TRAP) number."* — and the bed's row reads *"set up as INSTRUMENTAL (HIP-HOP / TRAP), fader −4.0 dB, already mastered - voicing kept light"*. When the vocal sang for under two seconds of the window, when the bed dropped out, or when the bed is locked or wasn't part of this run's setup, the act **declines and says so** rather than guess; the bed's fader is then yours. It always is afterwards, too: the ratio is a starting point in the fader's own units, nudge it and it stays.

An untyped, unnamed track next to a typed voice is typed an instrumental before it learns (the HIP-HOP bed under a RAP VOCAL, the POP bed under anything else) — only if its spectrum is broadband, so a still-unnamed kick is never mistaken for the bed. AUTO detection itself never declares a bed: a dense synth bus must not become one.

Sections — a stack comes down together

Six rhythm guitars that each landed on the guitar profile's level are not six guitars at the guitar's level — they are one guitar section eight dB louder than any of them, because the per-strip pass places every strip on its own and never asks who else is playing the same part. So after the tracks have settled, and before the 2-TRACK act, the run takes a census of the session by *family* — GUITAR, KEYS, SYNTH, PERC, ORCH, and the backing vocals, doubles and choirs behind a lead — and listens to every family with two or more members for ~6 s (longer with a longer listen length): **SECTIONS**. It measures a section the way you hear it: every member's real loudness, power-summed frame by frame, gated so a part that drops out for a bar doesn't count as a quiet part. The number that matters is the **stack** — how far the section sits above its loudest single member. The act then takes the whole stack off — every member comes down by it in **one common move, so the balance inside the section survives verbatim** — so however many guitars there are, they sum to one guitar's worth. The per-family level in the profile table is a *budget*, not a per-track offset, and the same holds for the backing vocals behind a lead: a lead plus its stack sums to the table's vocal level.

What the census knows that a sum doesn't: the lead vocal and the vocal bus are the anchor, never in a section; kick, snare, toms, overheads and rooms are *one kit* (only whole drum sources — a drum bus, a DRUM KIT, a top loop — stack, and only two or more of them); a bass and its sub are layering, not a stack; a lead synth or an orchestral melody is not part of its family's bed; FX never move. Tracks a bus already carries are left to the bus — its own calibration holds the sum, so its level is the section's level.

A section that takes turns instead of stacking (one guitar in the verse, another in the chorus), one that barely played, one that stacks under a dB, one that reads louder than that many parts can add up to, or one whose members are all locked is **left alone, and the report says why**. Strips that never played and locked strips are measured but never moved; only strips this run set up move at all (RE-BALANCE reports the stack and points you at a from-scratch run, so a second press never doubles the move). A move of 3 dB or more re-plans the carves before the run goes on. The report says it in numbers — *"Six guitars, playing together most of the time: they stacked to 7.6 dB over the loudest one, so the section came down 7.6 dB."* — and each member's row shows its fader move. Afterwards the faders are yours; one Undo takes back the whole run.

Stems — the five placed against each other

SECTIONS is pass one: inside a stem, like parts share one budget. **STEMS** is pass two. Once the tracks have settled (after SECTIONS and 2-TRACK, still in the track tier — so a bus then learns the placed mix), the run measures the five stems — **drums** (percussion included), **bass**, **vocals**, **instruments**, **FX** — by gated loudness, power-summed from every member frame by frame, and places each at its budget under a reference stem: the drums when their kit played — a kick, snare, clap, overheads, room mics or the whole kit — else the next stem in that order that played. The table: drums 0, bass −2, vocals −1, instruments −3, FX −12 LU under the reference. One trim per stem, applied to every member, so the balance pass one built inside a stem survives verbatim — **up as well as down**, because a budget is a place, not a ceiling. The reference stem never moves; BALANCE decides where the whole mix sits afterwards. A drums stem that is only percussion — a shaker, a hat loop, no kit — is measured and reported but never placed, and one pass never moves a stem more than 18 dB: a bigger number is a measurement problem, so that stem declines and the report says so.

What it leaves alone, and says so: a stem a bus carries (the bus's own level places it), a stem nothing in this run set up, a stem that barely played, a locked stem, a soloed run — and **every 2-track session**, where the vocal-to-bed ratio owns the shape. Any stem move re-plans the UNMASK carves once, on top of the finished balance: unmasking is the last thing that happens, because that is where it belongs. The report reads *"The instruments were 5.4 dB off their place, so the whole stem went back 5.4 dB: 2.0 LU under the vocals."*

Your moves survive a run — and a mix that sits still

Every card remembers where the last run left its fader and pan. **What you move after that is yours.** On the next run the strips re-learn with your offset on top of the new profile level; SECTIONS, 2-TRACK and STEMS measure *around* it, so they make the decisions they would have made on the untouched mix; SPREAD THE STAGE leaves a pan you placed where you put it (the report says *"2 pans you placed yourself stayed where you put them"*); and at the end the ledger is written back with your offset preserved, so it survives the run after that too. A strip's own LISTEN re-bases its line — that learn wrote the fader itself. The ledger persists with the session and rides in the mesh-wide undo.

And a second press does not re-litigate what the first one decided: each act remembers the trim it applied and the cast it applied it to, and a later run whose own measurement lands within 1.5 dB of that on the same cast replays the remembered number — the run-to-run spread of the same placement is under a dB of calibration lottery, and a fader that moves 0.6 dB "for no reason" is exactly what makes a mix feel like it will not sit still. Only a run that re-set the members up remembers or holds, so RE-BALANCE can neither replay nor overwrite it. A member added or gone, or material that genuinely changed, clears the hold by its own size. The report says *"Same cast as last time, so I kept the last run's numbers rather than re-deciding them."*

REVISIONS — the next ask, one click away

The revisions a producer asks for after a first pass are nearly always the same handful, and the run already knows what it measured. So when AUTO-MIX finishes, a **REVISIONS** band appears under the mixer header with up to four chips — only the ones this session earns, never a fixed menu. **VOCAL FORWARD** needs a lead vocal that played; **BACKING VOCALS BACK** needs support voices that played; **MORE BASS** needs a placed bass stem; **LESS ROOM** needs overheads or room mics that sit within 6 LU of the kick or the backbeat; **GUITARS WIDER** needs two guitars. Hover a chip and it tells you the evidence in the run's own words — *"LeadVox sits 3.6 LU under the drums; its 7 support voices sit 3.0 LU under it"* — and what one click does.

One click is one notch. VOCAL FORWARD brings the lead 1.5 dB forward and pulls every support voice 1 dB back, so the lead comes forward and still sits in the stack. BACKING VOCALS BACK moves them 2 dB as one move — the balance inside them is untouched. MORE BASS is 1.5 dB on the whole bass stem (the kick duck is already armed, so the low end keeps its order). LESS ROOM takes the room mics down 2 dB. GUITARS WIDER steps every guitar out one width class of the stage spread, on its own side. Three notches is a different mix, not a revision, so the fourth click takes the chip back to zero — click through and the cast is back where the run left it. **Shift-click walks the other way**; below zero the chip reads the other way too — VOCAL BACK, LESS BASS, MORE ROOM, GUITARS TIGHTER. The dots on the chip are the count. Each click speaks once, in the report's voice: *"I brought LeadVox 1.5 dB forward and pulled its 7 support voices 1 dB back, so it still sits in the stack."* One press of UNDO takes the click back, mesh-wide, like every other move. The AUTO-MIX COMPLETE report itself stays a report; the chips live under the mixer header only — on every instance's mixer, not just the one you pressed AUTO-MIX on, and a mixer you open later gets them too.

A revision is yours, and it survives the run. It is written to the mix ledger as your offset — never as a placement — so the next run re-learns the strips with it on top, measures around it, and lands it again untouched; the report then says *"You asked for the vocal forward (1.5 dB) and more bass (1.5 dB); both are kept"* and each ledger row carries *"your +1.5 dB kept"*. Clicks stack across chips per card (the backing vocals carry both VOCAL FORWARD's step and their own), a click on the same chip moves its count rather than adding a second move, and a fader you dragged in between is left where you put it. A locked strip is never touched by a chip and is named when it was skipped; a chip the session cannot offer is simply absent, and one a law blocks is dimmed with the reason. The offer is spent by the next press, which makes a fresh one — but a chip that still carries a count always comes back, whatever the new run measured, so you can always see it, cycle it or switch it off; the counts and the faders persist with the session. Pans placed by GUITARS WIDER are pans you placed, and a later run leaves them alone.

Locking a strip out of AUTO-MIX

Every card's type picker carries a footer toggle: **"LOCK FROM AUTO-MIX" → "LOCKED - AUTO-MIX KEEPS ITS HANDS OFF."** A locked strip is never set up, re-levelled, carved into, re-armed, panned or fader-moved by any run — including from-scratch runs, RE-BALANCE and PHASE ALIGN — and wears a small warm padlock on its type chip. It still *counts*: it stays in the masking picture, keeps its pecking rank (others still step back for it), and its level still informs the balance pass. Lock the strips you've hand-tuned, then let AUTO-MIX work around them.

The run

AUTO-MIX snapshots every instance first — **one Undo takes back the entire run, mesh-wide** — then dims your monitors 12 dB for the duration (released the moment it ends, however it ends) and works

tier by tier: tracks first, then group buses, then the master, so a bus always learns from tracks that have already stopped moving. Per tier: signal-gated learns → (tracks only) the UNMASK pass, the stage spread and the kick-duck arming → settle until every level window lands. The final act, **BALANCE**, measures the real sum and pulls every contributor down together if the session sits above the mix target.

Section memory keeps the levels honest. Every strip remembers the loudest section it has heard in the last minute of playback, and learns level against *that* — so a run pressed over a quiet verse still sets the trims the chorus needs, instead of a mix that explodes at the drop. When it happens, the report says so, per strip.

A bus learns the mix as it is now. A group bus or the 2-bus has been hearing the raw, un-set-up sum long before you press — and once the run has re-levelled the tracks feeding it, that memory would describe a mix that no longer exists. So when a run has touched the tier beneath a bus, the bus forgets its section memory right before its own learn and calibrates to the sum it actually receives. An incremental press that only adds a bus keeps the memory: nothing changed beneath it.

The level check waits for the strips to be still. The final pull-back only measures the sum once every strip's own level correction has finished settling — and if a bus is still calibrating when the check runs, the run leaves the overall level alone and says so in the report, rather than correcting a level that is about to change on its own. RE-BALANCE once it has landed.

You watch it narrate itself — "LISTENING TO YOUR TRACKS", "UNMASKING - MAKING SPACE", "SETTLING TRACK LEVELS", "LEVELLING THE SECTIONS" (the SECTIONS act, when a family has two or more members), "BALANCING THE VOCAL AGAINST ITS BED" (the 2-TRACK act, when there is a pair), "LISTENING TO YOUR MASTER", "BALANCING THE MIX LEVEL", amber "WAITING FOR PLAYBACK - PRESS PLAY" — and it **needs the session playing**: silence pauses it, it never guesses.

Stopping is safe: the button reads **STOP** mid-run, ESC does the same, and a stopped run says exactly where things stand: *"Everything it had already done is still in place, and the strips it never reached kept their settings. Undo takes back the whole run."*

The decision ledger

The result is a modal, not a toast — "AUTO-MIX COMPLETE". On top, the run tells you what it did in its own voice — the same voice as SMART SETUP's report, greeting you by name and walking the decisions in order: *"Hey Phil! The mix is in — I set up 3 strips and placed them in your order... I read the signal types and put LEAD VOCAL in front — everything else steps back for it... 9 carves keep the tracks out of each other's way — the deepest pulls BACKING L back 3.5 dB at 3.1 kHz for LEAD VOCAL... Then I pulled the whole session back -3.5 dB in one move."* The phrasing varies press to press; the facts never do. Under the prose sits the ledger — **one line per strip saying exactly what the run did to it** (it scrolls, so every line is reachable): *"set up as LEAD VOCAL, fader +2.3 dB, panned 45% L"*, *"re-levelled"*, *"kept its settings"*, *"stepped back for LEAD VOX"*, *"locked - untouched"*, *"never played - untouched"*. When a strip's learn was less than certain, the ledger says why in words — *"(heard a quiet section - level set for the loud one; short capture)"* — never a bare confidence number. Nothing the run did is a mystery, and it always ends: *"Undo takes back the whole run."*

A good session for AUTO-MIX: strips on your tracks (MINIs are fine — that's what they're for), a full strip on the 2-bus, signal types declared where you know them (or trust AUTO), hand-tuned strips locked, and the busiest section of the song looping. Then press the button and read the ledger.

PHASE ALIGN — one press, every mic in time

Multi-mic recordings — drum kits, guitar cabs close+room, top and bottom snare, DI plus amp — arrive at each mic at slightly different times, and sometimes flipped. The fix used to be zooming into waveforms and nudging regions. Here it's one button.

The press

ALIGN sits in the MIX view header, between UNDO and UNMASK. Press it while the session plays: every playing, unmuted strip records **two seconds of its raw input, locked to the host timeline** (soloed strips narrow the scope to themselves; AUTO-MIX-locked strips sit out). The button reads **LISTENING** while it captures — amber **PRESS PLAY** if the transport is parked — and the toast says what's happening: *"ALIGN — KEEP THE SESSION PLAYING; THE STRIPS ARE LISTENING TO HOW YOUR MICS ARRIVE."* It needs at least two playing strips — *"it lines up mics that hear the same source."*

What it does

The overhead is the anchor when you align a group that has one — it's the only mic in the room that hears every close mic, so it's the only honest reference for both polarity and time. (Against a kick mic, the snare, toms and hats read as "nothing in common with the anchor" and get skipped, and the polarity reference runs backwards: the overhead ends up flipped to meet the kick.) OH L and OH R break the tie on the pecking order. With no overhead in the group it's the kick, then the pecking order's #1, then the first candidate. Ride and cymbal spot mics never anchor. Aligning the whole mesh, or a solo'd set, is unchanged. Every other capture is cross-correlated against it, and where a **real relationship** exists, the strip gets exactly two things: a **polarity flip** if its mic is inverted, and a **sample-exact, delay-only time shift** (up to 20 ms). Tracks that share nothing with the anchor — your synth bus during a drum alignment — are **left alone**, stated in the receipts, never guessed at.

Everything is re-anchored on the **latest** arrival, so every shift is a delay and nothing ever plays early — no latency is reported, because a moved mic isn't latency, it's the fix. The shift lives outside the INPUT stage's enable (switching INPUT off can't slide you back out of phase), bypass honestly plays the un-aligned mic (that's your A/B), and re-running measures the *residual* — a mic you flipped by hand is honored, never fought.

PHASE ROTATE — the third thing a mic can be wrong by

Flipping polarity is one bit. A delay is one number. But two mics on one source are often out by something in between — the low end cancels, and no amount of sliding fixes it. That's what a **phase rotation** is for, and ALIGN now solves it as a second pass.

Where a pair still matches poorly after the flip and the slide, the hub tries phase corners in half-octave steps from 30 Hz to 960 Hz, re-solving the lag and the polarity at each one, and takes a rotation only when it removes **at least half** of what was left over — a turn you wouldn't hear is a turn it doesn't make. You get the result on the INPUT card's **ROTATE** well, and the tile says **ROT**. It reports no latency, and re-running ALIGN measures the residual and stays put, exactly like the delay and the flip already did.

One case it won't act on, and says so out loud. A rotation of this kind can only ever add lag to the strip it sits on, so when the mic that needs turning is the *anchor*, nothing done to the other mic can fix it. ALIGN

never rotates the anchor behind your back. The receipts tell you instead: *"turns the other way: ROTATE the anchor (OH) at ~85 Hz by hand"* — one well, one drag, and you're there. Far mics — anything typed as a room mic, or more than 6 ms out on the flat arrival — are never rotated at all; distance is the sound, not a fault.

The receipts

The result is a modal — **"PHASE ALIGN COMPLETE"** — with one line per strip, with the measurement: *"SNARE BOT: flipped, moved 3.9 ms - match 0.31 → 0.88"*, *"OH L: already in phase"*, *"KICK IN: anchor"*, *"SYNTHS: nothing in common with the anchor - untouched"*. The prose sums it up — *"Align done - 3 strips moved into time, 1 polarity flipped, re-anchored on the latest arrival so nothing plays early"* — and, as always: *"Undo takes back the whole run."* One mesh-wide undo step; a cancelled run (STOP or ESC) applies nothing.

Pro tip: Solo the mics that belong together first — the kit, or the two guitar mics — so ALIGN works on exactly that family, with your most important mic ranked #1 so it anchors.

THE_STRIP 3 MINI — the companion

What it is

THE_STRIP 3 MINI is **the same engine as the full strip**, narrowed to the five stages that matter on the tracks you have forty of: **INPUT, LOW-CUT, EQ, COMP and DUCK** — at a fraction of the CPU and with **zero latency, structurally**: no lookahead, no oversampling, nothing that could ever report a sample of PDC. It's safe on live inputs, tracking paths and forty-track deployments by construction, not by configuration.

It is not a lite product bolted on the side: parameter names, ranges, defaults and state are identical to the full strip, so a full-strip preset loaded into a MINI applies the five stages it has and ignores the rest — and a MINI channel drilled into from the full strip shows exactly the five stages it runs.

The face

One 620×620 window: a **hero XY pad** filling almost everything, and a slim module row underneath.

- **The pad** — vertical = output gain (±12 dB), horizontal = pan. Relative pickup, SHIFT fine, double-click reset; **LOCK PAN** and **LOCK GAIN** pin either axis (pan lock ships engaged — a MINI is a level first). Under the puck, a particle field breathes with the final output's stereo image.
- **DIM** — beside the gain / pan readout on the pad: the same temporary monitor cut (−12 dB) as the full strip's DIM. Lit while engaged; the readout keeps showing your gain, not the cut. Stays engaged across a reload until you switch it off.
- **THE CARVE** — the backdrop behind the pad: the live response of the MINI's six-band dynamic EQ, drawn as one faded path, dynamic movement included. You never edit these bands by hand — they belong to SMART SETUP and to the network. When UNMASK asks this track to step back for the vocal, **the carve is what you're looking at**. Empty state: *"NO CARVES - THE MIX HASN'T ASKED FOR ROOM YET."*
- **The module row** — three cards, each an LED plus one draggable headline value: **LOW-CUT** (frequency, 20–500 Hz), **COMP** (threshold — the same calibrated/REDUCTION semantics as the full strip), **DUCK** (amount, 0–100 %). Same gesture law as everywhere: drag, SHIFT fine, double-click reset.
- **SMART** — the last chip on the row: the full SMART SETUP panel, nearly full-window — the complete profile grid, YOUR PROFILES, CREATE PROFILE, 3/6/12 s LISTEN, everything. On a MINI the learn configures the five stages it runs. During capture the chip is a progress ring; "WAIT" while it's starved.
- The usual chrome: bypass (with the full-screen "BYPASSED" overlay), settings, docs, AI assistant, tooltip bar, update banner. One license covers both plugins — MINI registers as THE_STRIP 3.

There is deliberately **no mixer in the MINI** — no STRIP|MIX|HYBRID toggle. A companion you put on forty tracks should not offer forty doors into the same mixer; the mixer is the hub's job.

A full citizen of the mesh

Drop MINIs on your tracks and a full THE_STRIP 3 anywhere in the session, and from the full strip's MIX view every MINI is a complete channel: fader, pan, mute, solo, groups, loudness/TP/correlation readouts (with the "MINI" badge so you know what you're looking at). You can:

- **Learn it remotely** — alt-click its type chip; the capture ring runs on the card, the MINI's window never has to open.
- **Drill into it** — EDIT STRIP shows its five stages, live visualisations from its audio.
- **Carve it** — UNMASK writes its network bands like any track's; the MINI shows the result as THE CARVE.
- **Duck it** — its DUCK stage follows the session's kick over the mesh, and AUTO-MIX arms it like any bass-family track.
- **AUTO-MIX it** — full participant, including the mesh-wide undo.

Without a full strip in the session, everything local still works — the five stages, the pad, SMART SETUP, bypass. The carve stays empty (nothing is asking for room), FOLLOW KICK finds no kick and simply never fires, and there's no mixer view. MINIs alone still mesh with each other; they're just waiting for a hub.

The intended session

MINI on everything, THE_STRIP 3 on the vocal, the drum bus and the 2-bus. Type the tracks (or let the names do it), loop the loudest section, press **AUTO-MIX** on the hub. From then on you mix from the hub's MIX view — faders, the pecking order, the ducks — and open a MINI only to ride its pad.

Visual Feedback & Metering

- **Every card is a meter.** Live per-module visualisations in each rack card's well: input meters, GR curves with numeric readouts, the dynamic EQ composite, goniometer, harmonic bars, echo taps.
- **The gaps are meters too** — per-stage level, GR pip, and the measured in→out delta at every module boundary. The internal-headroom **HOT** warning fires from these taps.
- **The gain map** on the OUTPUT page is the whole story in one line: in-level, every stage's contribution, the strip's own correction, out-level.
- **Loudness suite** — momentary / short-term / integrated LUFS, range, held true peak, correlation; -14 LUFS horizon on the meter; click-to-reset here and on every mixer card.
- **The steel breathing ring** is the calibration's tell: anything it is about to rewrite breathes first. You will never be surprised by a self-set knob.
- **Editors draw the engine's math** — the low-cut curve, the duck figure, the MOD LFO (with a playhead on the real phase), the comp transfer curve: all rendered from the same code that processes the audio, so picture and sound cannot disagree.
- **Bypassed stages draw honest flat traces** — a scope never freezes (frozen reads as crashed), it shows the truth: nothing happening.
- **Mixer cards**: stereo meters with scale ticks, standing dB readouts, pip rows, heat ribbons, type chips that double as learn-progress bars, and status lines that speak one vocabulary.
- **Tooltips in five languages** (EN/DE/FR/ES/IT) on every zone of both plugins, in the bar at the bottom.

How to Use It — Real-World Workflows

The 60-second vocal chain

1. Insert THE_STRIP 3 on the vocal. It opens pre-typed from the track name.
2. Loop the loudest phrase, press **SMART**, confirm LEAD VOCAL, hit **LISTEN**.
3. Read the report — it tells you what it did, in dB — then taste-tweak: COMP threshold (it's in dB of reduction), DE-ESS range, REVERB ducking.
4. **BYPASS** on the OUTPUT rail to A/B. MATCH is on, so you're comparing tone, not loudness.

Mix a whole session in an afternoon

1. MINI on every track, THE_STRIP 3 on the vocal, drum bus and 2-bus.
2. Declare types where you know them (the track names have usually done it).
3. Loop the densest section. **AUTO-MIX** → answer the preflight (WET? / FOUR TO THE FLOOR? / SPREAD THE STAGE? / listen length) → RUN, then read the ledger. The SECTIONS line says how far each stacked family sat over its loudest member and what came down together.
4. When the modal reports, listen. Then adjust the *decisions*, not fifty knobs: drag the pecking order, ride faders in MIX view, promote whoever should win with their UNMASK pill.
5. Anything you hate: **UNDO** — the whole run, every instance.

Mix a vocal over a beat (2-track)

1. THE_STRIP 3 on the vocal, a strip (or MINI) on the instrumental. Name the instrumental "Beat" / "Inst" / "Playback" — or pick an **INSTRUMENTAL** genre in the 2-TRACK category — and the vocal opens as LEAD VOCAL or RAP VOCAL from its name.
2. Loop a verse where the vocal actually sings. **AUTO-MIX** → RUN.
3. Read the 2-TRACK line in the ledger: it tells you how far it moved the bed and where the vocal now sits (in LU), and whether the beat was left alone because it arrived mastered.
4. Too forward / too buried? Ride the bed's fader in MIX view — the ratio was its starting point, in dB, and nothing re-fights you. The carves keep following the voice either way.
5. Everything you hate: **UNDO** — the run, both strips.

The vocal that always cuts through

1. Give the vocal rank #1 (press UNMASK on its card, or drag it to the top of the board).
2. Guitars, keys and pads now carve **only while the vocal sings** — check their cards: "STEPS BACK FOR LEAD VOX".
3. Set **HOW HARD** to taste; **COMPARE** (hold it, or click it to keep it off) to hear the mix without the feature.
4. Lock anything that must never bend — the 808, the hook sample.

Kick and bass, solved

1. Type the kick track KICK DRUM and the bass track BASS (SMART SETUP or the mixer type chip).
2. On the bass card, press **DUCK** — "DUCKING UNDER KICK". Or run AUTO-MIX, which arms it for you.
3. In the bass's DUCK editor, set **BELOW** to ~150 Hz so only the lows step aside, and CURVE LEN to taste.
4. No kick in the arrangement (four-to-the-floor drop section)? The preflight's FOUR TO THE FLOOR? answer — or the duck's SYNC trigger — pumps on the grid instead.

Master polish

1. THE_STRIP 3 on the 2-bus, profile MASTER (or MASTER LOUD).
2. LISTEN writes the glue comp, the tonal touches and the limiter.
3. Set **CEILING** to your delivery spec, switch **TP** on, watch TRUE PK on the OUTPUT page.
4. Click the loudness column to reset, play the whole song, read INT and RANGE.

Surgical rescue on a problem source

1. GATE with TRIGGER = Transient for the bleed; DETECT BAND so it only hears what it should.
2. Dynamic EQ bells on the resonances — LISTEN per band to find them, DYN so they only duck when ringing.
3. DE-ESS in Split mode; LISTEN to what it removes until you hear no vowels.
4. Delta-solo each stage (**S**) to verify each one earns its place.

Honest A/B, always

- Per stage: the **S** delta-solo plays what that module contributes.
- Per strip: **BYPASS** + **MATCH** on the OUTPUT rail.
- Per session: **BYPASS ALL** in MIX view.
- Per feature: **COMPARE** on the UNMASK board.

Pro Tips & Creative Techniques

- **Calibrate first, think later.** One CALIBRATE (or any LISTEN) and every threshold in the strip becomes a real number. The strip is *designed* to be driven by the dB readouts.
- **REDUCTION mode is the whole compressor.** Decide how many dB of compression the source needs, set that number, done. Ratio becomes a flavour choice, not a math problem.
- **Reorder like you mean it.** De-esser *before* the compressor tames esses before they pump the GR; *after*, it cleans what compression exaggerated. Drag it, hear it, keep the better one.
- **The wet duckers are the modern move.** REVERB and DELAY DUCKING at 0.3–0.6 gives you wetter mixes that stay clearer than dry ones. SMART SETUP seeds it; MAKE IT WET forces it.
- **UNMASK replaces a rack of sidechain routing.** The old way was six compressors and a spider web of sends. Here it's one drag on one board — and it releases the moment the winner stops playing.
- **Promote, don't EQ.** When something disappears in the chorus, don't reach for its EQ — press its UNMASK pill. The mix steps aside instead of the track changing character.
- **BELOW on every duck.** Full-band ducking is for effect; ducking below 150 Hz is for mixing. Same knob, different job.
- **Trim lanes are your gain-staging conscience.** The delta readout between two cards tells you what a stage *actually* did to the level. If a stage's delta and its musical benefit don't match, fix the trim, not the threshold.
- **MONITOR solos before commitment.** TRANSIENTS' Transient/Sustain monitors, band solos in MB-COMP, per-band LISTEN in the EQ, SC LISTEN on gate/comp/duck — the strip lets you hear every detector. Use them, then decide.
- **Delta-solo the WARMTH.** The S chip on WARMTH plays only the added colour. If the delta sounds like fizz, change the model; if it sounds like glue, push the drive.
- **MINIs everywhere is cheaper than you think.** Zero latency and a fraction of the CPU per instance — and the hub can learn, carve, duck and balance all of them without you opening a single window.
- **LINK the kit.** Your drum tracks grouped themselves the moment they were typed — collapse the group, ride one fader for the whole kit, and switch LINK on so the low cut and comp envelope stay identical across every mic. One knob, whole kit.
- **Cmd/Ctrl+Z is fearless.** SMART SETUP, AUTO-MIX (mesh-wide), UNMASK order changes, every knob, every reorder — all undoable. Experiment freely.

Keyboard Shortcuts & Gestures

Key / gesture	Action
TAB / Shift-TAB (also → / ←)	Walk the chain in signal order, land on OUTPUT after the last module, wrap
S	SMART SETUP panel
Type while SMART SETUP / the type picker is open	Search the grid; RETURN picks the first result, ESC clears the search, then closes
M	STRIP ↔ MIX view
B	Bypass the selected module
ESC	One layer per press: close the topmost overlay → stop a running AUTO-MIX → close the drilled channel → MIX back to STRIP
Drag a knob	140 px = full range; **SHIFT** = fine (480 px)
Double-click a knob / fader / pad	Reset to default / unity / centre
Right-click a knob / fader / trim lane	Type a value — "1.52k", "12.5 kHz", "-6 dB", "30%", "-inf" all parse
Wheel	Adjust knobs; Q on an EQ node; scroll the mixer
Drag a rack card	Reorder the chain (INPUT and LIMITER stay pinned)
Click a card LED	Stage on/off
S chip on a card	Delta-solo that stage
Double-click empty EQ curve	Add a band (double-click a node removes it)
Alt / right-click a mixer type chip	Run LISTEN on that channel, headless
Alt/Cmd/Ctrl-click any mixer S chip	Clear every solo
Alt-click AUTO-MIX	The from-scratch run
Double-click a mixer card name (local)	Rename the instance, session-wide
Click a loudness readout (OUTPUT page or mixer card)	Reset integrated LUFS + true-peak hold
Hold COMPARE (UNMASK board)	Hear the mix without UNMASK — a quick click latches it off until you click again
Drag a track (MIX MAP, armed)	Across = pan, towards you = louder; the first travel locks the axis, **ALT** frees both
Double-click a track (MIX MAP, armed)	Fader to 0 dB
Drag / double-click empty room (MIX MAP)	Turn the stage / reset the camera

Corner grip	Zoom the whole UI

Unused keys pass through to the host; while a modal is up, only ESC works.

Quick Start

1. **Insert** THE_STRIP 3 on a track. It has already guessed the source from the track name.
2. **Press SMART → LISTEN** on the loudest section. Read what it did; tweak what you'd change.
3. **A/B honestly** — BYPASS on the OUTPUT rail (MATCH keeps the loudness fair).
4. **Add more instances** — MINIs on the small stuff — and press **M**: the whole session is now a mixer.
5. **Press UNMASK** on the track that matters most. Drag the board until the order is your arrangement.
6. **Press AUTO-MIX** when you'd rather start from a mixed session than a blank one. Undo if you wouldn't.

Specs

THE_STRIP 3

- **Plugin type:** Channel strip / cross-instance mix system (audio effect with sidechain input).
- **Modules:** 15 — INPUT, LOW-CUT, GATE, EQ (6 dynamic bands, M/S/L/R per band), COMP, DE-ESS, DUCK, TRANSIENTS, MB-COMP (3 bands), STEREO, REVERB, DELAY, MOD (5 modes), WARMTH (6 models), LIMITER — fused in one engine; 13 freely reorderable, INPUT pinned first, LIMITER pinned last. Per-module enable (click-free crossfades), output trim, delta-solo on 13 stages (all but INPUT and MOD). The LIMITER carries a built-in anti-aliased CLIP stage (0–100 %, zero latency, bit-inert at 0).
- **Calibration:** -12 dB (configurable -24...-3) operating reference; window-based measure-and-hold auto gain-staging (input / per-stage / final); calibrated threshold readouts; compressor REDUCTION mode (-N = N dB of gain reduction).
- **SMART SETUP:** 81 factory profiles (incl. VOCAL DOUBLE, DRUM KIT, the PERCUSSION/KEYS/BASS/FX family buses, TOMS / ROOM MICS / RIDE, BRASS & STRINGS, ORGAN, MALLETS, VOCAL CHOPS, CHOIR, and six INSTRUMENTAL 2-track beds whose voicing adapts to how mastered the material already is) + user profiles from reference audio (1–8 files, WAV/AIFF/FLAC/MP3/OGG, optional online description matching); AUTO source detection; signal-gated 3/6/12 s capture with 60 s section memory (levels set for the loudest section heard, not the quiet bar you pressed on); MAKE IT WET; full plain-language report; one-step undo.
- **Mesh:** automatic cross-instance discovery (THE_STRIP 3 + MINI, one network). MIX view (faders ±60...+12 dB, pan, mute, session-wide bidirectional routing-aware solo, self-seeding family groups with collapse / group M-S / shared VCA fader / LINK-ganged low-cut + comp times + duck length, rename, per-card loudness/TP/correlation with click-reset, BYPASS ALL as a whole-session A/B — processing, faders, pans and trims out, mutes and solos held, one frozen shared sum offset), HYBRID rail view, remote drill-in editing with remote visualisations, headless remote LISTEN, bus/master auto-detection with manual override, bus family inheritance from carried tracks, per-bus "N MAKING ROOM BELOW" carve tally.
- **UNMASK:** draggable pecking order; pan-aware, perceptually weighted collision analysis; measured-width precision carves (wide bell on a smear, notch on a honk, depth sized by what's actually masked); at least 3 free bands guaranteed per SMART-SETUP'd strip, up to 4 carves; per-track lock; HOW HARD depth; mesh-wide COMPARE (hold or latch, a crossfaded band suspend that never reaches automation or the session); kick/bass pocket handled by duck, not EQ; state persists with the session.
- **AUTO-MIX:** incremental by design (full setup only where needed, level-only elsewhere, RE-BALANCE when nothing is new); 4-question preflight (MAKE IT WET? / FOUR TO THE FLOOR? / SPREAD THE STAGE? / listen length 3-6-12 s); symmetric stage spread with the human vocal rule (one centred lead, extra leads go wide); per-strip AUTO-MIX locks; tiered runs (tracks → buses → master) with 12 dB monitor dim and section-memory-honest levels; 2-TRACK act for a vocal + bed pair (gated momentary-loudness ratio measured over ~6 s, target derived from the two profiles, bed fader only, declines rather than guesses); SECTIONS act for stacked families (type census as the prior, a gated ~6 s power-summed measurement as the veto, stack = section minus its loudest member, one common trim of the whole stack — leads, kit pieces, bass, bussed tracks and FX never

moved, declines rather than guesses); STEMS act (pass 2: drums / bass / vocals / instruments / FX by gated loudness at 0 / −2 / −1 / −3 / −12 LU under the drums or the next reference that played, one trim per stem, at most 18 dB per pass, percussion-only drums never placed, up as well as down, carried stems and 2-track sessions decline, then one carve re-plan on the finished balance); the mix ledger (fader and pan moves made after a run survive the next one as offsets the acts measure around); the re-press hold (the same cast landing within 1.5 dB of last time replays last time's trims); cancellable (STOP/ESC); per-strip decision-ledger result modal; one mesh-wide undo step.

- **PHASE ALIGN:** one-press multi-mic alignment from the MIX header — 2 s transport-locked captures, cross-correlation against the pecking-order anchor, polarity flips + sample-exact delay-only shifts (≤ 20 ms) only where a real relationship exists, re-anchored on the latest arrival (nothing plays early, no reported latency), per-strip receipts with before/after match figures, one mesh-wide undo step.
- **Sidechain:** host sidechain bus feeds GATE, COMP and DUCK; the mesh feeds DUCK (Follow Kick) and EQ (UNMASK) without routing.
- **Oversampling:** WARMTH (default 2×) and LIMITER (default 1×) independently 1–8×; global quality switch FAST / CLEAN (CLEAN adds a guard-filtered decimator for a far lower alias floor at higher latency).
- **Limiter ceiling:** sample-accurate by default; optional TP true-peak stage (1/8-sample ISP detection, ~2 ms) makes it a true-peak ceiling at any OS setting.
- **Latency:** reported honestly and only for what's engaged — lookaheads (gate ≤20 ms, comp ≤5 ms, duck ≤20 ms, limiter ≤20 ms, default 1.5 ms), oversampling filters, TP (~2 ms). Master bypass reports zero. Reported = measured.
- **Metering:** per-stage boundary meters + GR, loudness suite (M/S/I LUFS, LRA, held dBTP, correlation, -14 LUFS horizon), gain map, spectra per module, goniometer, internal-headroom HOT warning.
- **UI:** 1200×720, UI zoom, full undo/redo (parked/restored across remote drill-ins), 5-language tooltips (EN/DE/FR/ES/IT), in-app documentation, AI assistant, update banner.
- **Formats:** VST3, AU (macOS), AAX — macOS (Apple Silicon & Intel) and Windows.

THE_STRIP 3 MINI

- **Plugin type:** Companion channel strip (audio effect with sidechain input), same engine as THE_STRIP 3, lean profile.
- **Stages:** INPUT, LOW-CUT, EQ (network/SMART-owned — shown as THE CARVE), COMP, DUCK. Identical parameters, ranges and state format to the full strip.
- **Latency: zero, structurally** — no lookahead, no oversampling, no TP; nothing can make a MINI report PDC.
- **UI:** 620×620 — hero XY pad (±12 dB gain / pan, axis locks, goniometer field, live carve backdrop), three one-knob module cards, SMART chip with the full SMART SETUP panel; 5-language tooltips, docs, AI assistant.
- **Mesh:** full citizen — mixed, soloed, grouped, metered, drilled into, learned remotely, carved by UNMASK, ducked by the session's kick, included in AUTO-MIX and its mesh-wide undo. "MINI" badge on its mixer card. No mixer view of its own by design.
- **Licensing:** one THE_STRIP 3 license covers both plugins.
- **Formats:** VST3, AU (macOS), AAX — macOS (Apple Silicon & Intel) and Windows.

Stuck, or found something this manual doesn't cover? Write to me at philspeiser.com/support. I answer support myself.