MENTESH2A · USER MANUAL

MENTESH2A — USER MANUAL

Leveling Amplifier · Model MENTESH2A · MENTESH ENGINEERING · v1.0

> *Looking for the 1-page quick-start? See* **HIGHLIGHTS.pdf**, *in the same folder as this manual.*

---

1. WELCOME

MENTESH2A is a faithful software emulation of the Teletronix LA-2A optical leveling amplifier, modeled against a specific 2014 Universal Audio LA-2A reissue (drawing C-10953 rev C topology). Every stage of the original audio path — input transformer T1, T4B optical attenuator, the 12AX7A / 12BH7A tube cascade, negative-feedback loop, and Sowter 8940 output transformer — is reproduced from schematic and verified by direct measurement. The plug-in is calibrated to a unit-capture session (140 program-stimulus captures, 2026-04-29) and tracks the hardware to within ±2 dB of measured drift across its operating range. It is built to be lived-in, dropped on a vocal, and forgotten about — exactly like the box it descends from.

---

2. INSTALLATION

2.1 MACOS (PKG INSTALLER — RECOMMENDED)

The shipped `MENTESH2A-1.0.0.pkg` is a standard macOS installer. It installs everything in one go — VST3, AU, CLAP, the Standalone app, the manual, and the factory presets — to the system locations every DAW already scans. No dragging, no manual folder navigation.

1. Quit your DAW.
2. Double-click `MENTESH2A-1.0.0.pkg`.

3. Click through the installer. Authenticate with your Mac password when prompted (admin rights needed because the targets are under `/Library/`).
4. Re-open your DAW. On first launch the host rescans and picks up MENTESH2A under the **MENTESH ENGINEERING** group.

The installer drops files at:

```
/Library/Audio/Plug-Ins/VST3/MENTESH2A.vst3
/Library/Audio/Plug-Ins/Components/MENTESH2A.component       (Audio Unit)
/Library/Audio/Plug-Ins/CLAP/MENTESH2A.clap
/Applications/MENTESH2A.app                                  (Standalone)
/Library/Application Support/MENTESH ENGINEERING/MENTESH2A/
    Manual.pdf
    Highlights.pdf
    Presets/Factory/*.json                                   (8 factory presets)
```

If you want only some formats, click **Customize** in the installer and uncheck the rows you don't need.

2.2 GATEKEEPER NOTE (UNSIGNED BETA BUILDS)

The beta installers are not yet Apple-notarized. macOS will warn:

> *"MENTESH2A-1.0.0.pkg" can't be opened because Apple cannot check it for malicious software.*

To bypass: **right-click** the `.pkg` file → **Open** → confirm in the dialog. This adds a one-time exception. You won't see the warning on subsequent re-opens or upgrades.

If a host complains the plug-in is "damaged" after install, clear the quarantine flag:

```
sudo xattr -dr com.apple.quarantine /Library/Audio/Plug-Ins/VST3/MENTESH2A.vst3
sudo xattr -dr com.apple.quarantine /Library/Audio/Plug-Ins/Components/
MENTESH2A.component
sudo xattr -dr com.apple.quarantine /Library/Audio/Plug-Ins/CLAP/MENTESH2A.clap
```

then rescan in the DAW. Production releases will be notarized and won't hit this dialog.

2.3 MANUAL INSTALL (ADVANCED)

If you'd rather drop individual binaries (e.g. installing for one user only, or in a VM): the same `.vst3` / `.component` / `.clap` are available on request as a separate zip. Drop them into:

```
~/Library/Audio/Plug-Ins/VST3/
~/Library/Audio/Plug-Ins/Components/
~/Library/Audio/Plug-Ins/CLAP/
```

(per-user paths — note the `~`). The PKG always installs to `/Library/` (system-wide); the per-user dirs are only used for manual drops.

2.4 UNINSTALL

Run from Terminal:

```
sudo rm -rf /Library/Audio/Plug-Ins/VST3/MENTESH2A.vst3
sudo rm -rf /Library/Audio/Plug-Ins/Components/MENTESH2A.component
sudo rm -rf /Library/Audio/Plug-Ins/CLAP/MENTESH2A.clap
sudo rm -rf /Applications/MENTESH2A.app
sudo rm -rf "/Library/Application Support/MENTESH ENGINEERING/MENTESH2A"
```

Your User presets and license are preserved at `~/Library/Application Support/MENTESH ENGINEERING/MENTESH2A/Presets/User/` and `~/Library/MENTESH ENGINEERING/MENTESH2A.lic` respectively. Delete those manually if you want a clean wipe.

---

3. FRONT-PANEL CONTROLS

The faceplate is laid out left-to-right roughly the way a hardware LA-2A is, with a few additions plumbed in from MENTESH ENGINEERING's extended spec.

3.1 PEAK REDUCTION (PR)

Large rotary knob, right of the VU. Range **0 to 110**, default **40**. Printed `0`, `10`, `20`, `30`, `40`, `50`, `60`, `70`, `80`, `90`, `100`, `MAX`. The `MAX` sector past 100 is the "OAX" region — same control, extra travel for cell saturation and sidechain over-drive territory that the original front-panel scale stops short of.

PR sets **sidechain pre-gain** — i.e. how hard the detector is driven. It is **not** a threshold knob: threshold is fixed by the V3 + V4 bias points in the circuit. Higher PR pushes more signal into a fixed threshold, so you get a lower effective threshold and more compression.

3.2 GAIN

Large rotary knob, left of the VU. Range **0 to 100**, default **40**. Post-compression makeup level into the modeled tube cascade.

Crucially, GAIN does **not** affect the amount of compression. The sidechain taps off the T4B output node, before the GAIN attenuator. So you can set PR for the desired ride and then use GAIN purely for output trim, exactly as on the hardware.

3.3 LIMIT / COMP TOGGLE

Bat-handle switch in the bottom-left cluster. **COMP** (down) is the default. **LIMIT** is the harder mode.

On the hardware, SW3 toggles the 2.7 kΩ resistor R7 in or out of the T4B → V1A grid coupling:

- **COMP** — R7 in series. Softer coupling, slower effective onset, gentler ratio.
- **LIMIT** — R7 bypassed. Photocell drives the tube grid directly. Harder, more abrupt engagement, higher effective ratio.

The sidechain math is identical in both modes — only the audio-path coupling changes. This matches the schematic: there is **no sidechain time-constant change** when you flip COMP/LIMIT.

3.4 PHASE INVERT

Bat-handle switch with the Ø symbol above it, next to LIMIT/COMP. Inverts output polarity 180° (flips the sign of the sample stream after the full DSP chain). Useful for null tests and for fixing polarity flips in tracking.

3.5 MIX FADER

Small vertical fader in the bottom-left cluster. Range **0 % to 100 %**, default **100 %**. Standard wet/dry blend at the plug-in output: 100 % is fully compressed, 0 % is dry signal only. Intermediate positions give parallel compression. The numeric percentage is printed under the fader.

3.6 SATURATION INSERT

Horizontal strip across the bottom of the faceplate, below the VU. This is a separate insert — independent of the LA-2A topology — patched into the signal chain post-compressor and pre-makeup-gain. It draws from the same family of soft-saturation circuits MENTESH ENGINEERING uses elsewhere.

- **DRIVE slider** — horizontal fader on the left of the strip. Range **0 to 100** (printed as a 0-100 readout in the upper right). Sets pre-gain into the saturator. Models a tube-style curve with mild asymmetric clipping for 2nd-harmonic warmth (the dbx-162 *STOP LEVEL* character — louder + grittier as you push it). Default **0** (off — the insert is bit-transparent until you raise it).
- **LOW / MID / HIGH** band buttons — small lit pads to the right of the slider. Each one toggles whether that frequency band gets fed into the saturator. Default: **MID on**, low and high off (vocal- forward starting point).

- **LIFT / NOTCH** mode toggle — switches between two opposite-direction voicings. **LIFT** (orange) saturates the enabled band through the tube curve — adds energy + harmonic content (warmer, brighter, more apparent loudness). **NOTCH** (blue) does NOT saturate — it linearly attenuates the enabled band by the Drive amount (Drive=100 = band fully muted). Reductive shaping, the opposite of LIFT.
- **MATCH** button — when LIFT is engaged, it adds level. MATCH compensates the saturator's output back to roughly the level of the clean signal so you can A/B by ear without the louder version winning automatically. In NOTCH mode MATCH is a no-op (NOTCH already reduces energy).

3.7 VU METER AND METER SELECTOR

The center VU has a small chrome rotary above it labelled **METER** with three positions:

- **GR** — gain reduction. Needle reads negative dB of compression applied (push to the left).
- **+4** — output level referenced to +4 dBu nominal (pro line level).
- **+10** — output level reference for hotter / consumer-level metering.

The meter ballistics are standard 300 ms VU integration, driven from real DSP RMS pushed at 30 Hz from the audio engine.

3.8 POWER

Bat-handle toggle in the far-right cluster, labelled `ON / POWER`. Default **on**. When off, the plug-in fully bypasses (the host sees input → output unchanged, with the VU dimmed and the tube glow gone). This is the standard host-bypass route — it is the same thing as clicking the host's bypass button.

3.9 SCF — SIDECHAIN HPF

Small bakelite knob in the right cluster, labelled **SCF**. Range **0 Hz to 400 Hz**. Default **0 Hz** = filter **off** (the entire SC HPF is bypassed at 0).

Inserted in the sidechain detector path, it high-passes the signal feeding the envelope follower so low-frequency content (kick drum, sub-bass program material) doesn't trigger gain reduction. Useful on the mix bus and on a bass-heavy vocal where low-end thumps would otherwise dominate the ride. The current Hz readout is printed under the knob.

> *Why is the default off? Because LF-sweep measurement of the actual hardware (2026-04-22 capture session) found the LA-2A's sidechain is essentially flat in steady state — the famous "bass punch" comes from temporal envelope behavior, not a frequency-domain rolloff. SCF lets you add the rolloff yourself when modern program demands it.*

3.10 EMPH — EMPHASIS SCREW-TRIM

Smaller screw-trim style control, right cluster below SCF. Models the **R37 LIM/RESP** rear-panel trim pot on the hardware. Range **0 % (factory-flat) to 100 %**. Default **0 %**.

At 0 % the sidechain shelf sits at the as-shipped voicing of a kSilver- fitted unit (mild, audibly flat). Increasing toward 100 % adds an HF shelf (centered around 5654 Hz, fitted from the hardware's measured shelf response, +6.9 dB at the fitted center; up to roughly +12 dB at the top of the control's range). Use it for **de-essing** — high- frequency program drives the sidechain harder, so the compressor ducks sibilants without you having to insert a separate de-esser.

3.11 STEREO LINK (∞)

Small dot under the VU window with an `∞` symbol engraved below it. Acts as the stereo-link indicator and the click target. Tap the dot to toggle.

When OFF (dim), the plug-in operates as two independent mono channels — left and right have their own envelope followers and own GR. When ON (orange glow), the channels share a sidechain (envelope summed across L+R) and apply matched GR to both — so a transient on the left ducks the right by the same amount, preserving stereo image. The link knob in APVTS goes 0–100 % (**defaults to 0 % — unlinked at boot**); clicking the front-panel ∞ button toggles the link between OFF (0 %) and **engaged at 85 %** (not 100 % — see note below).

> **Default OFF at boot (since beta5).** *Earlier betas defaulted to ON. The boot default flipped to OFF so that engaging stereo link is always a conscious choice, matching the bus-comp convention of SSL / UAD / Plugin Alliance. On stereo program material* **you almost always want to enable it** *— see §6 below for the when-to-use-which table.*

> **Why does ∞ engage to 85 %, not 100 %?** *The real LA-2A is mono. A stereo setup means two physical units side-by-side, which never link perfectly — slight unit-to-unit drift always leaves a small amount of independence between channels.* **85 % models that real-world two-unit coupling.** *Power users who want hard-link can still drag the continuous link slider all the way to 100 %; the 0–100 % range is unchanged. The 85 % engage value is operator-locked per the LA-2A roadmap (see Phase 11.1 commit* `72c22fd`*).*

3.12 FACEPLATE VARIANTS

Open via the ⚙ **Tweaks** button in the top-right toolbar. The first control in the panel is **Faceplate** — a dropdown with five entries.

All five faces share the same compressor architecture, the same gain calibration (`kTubeChainDb`), the same saturation engine, and the same tube/transformer chain. The per-face variance is **two things only**:

1. **Photocell time-constant scale** — the envelope follower's attack and release tau values are multiplied by a per-face factor, modeling the documented ±30 % unit-to-unit T4B opto-cell spread (UA / Teletronix service notes).
2. **Output HF / LF tilt** — a one-pole high-shelf and one-pole low-shelf on the audio output stage (post-tube, post-T2), modeling per-cell HF response variance plus transformer / tube selection differences.

FACE	ATTACK	RELEASE TAIL	HF TILT	LF TILT	CHARACTER
Silver Teletronix *(your unit)*	10 ms	60 ms + 2.5 s	—	—	The reference. The UA in our studio (factory Sowter 8940 OPT + RCA NOS tube roll = bright sweet). Untouched.
Grey (UA era)	12 ms	60 ms + 4 s	−2 dB @ 6 kHz	—	Stock UA reissue with stock current-production tubes. Slower release, darker top.
MENTESH Signature (Black)	7 ms	50 ms + 1.2 s	+1 dB @ 8 kHz	—	House aggressive. Fastest cell. Modern bright punch.
Deep Purple	15 ms	80 ms + 5 s	−3 dB @ 5 kHz	−2 dB @ 60 Hz	Slowest cell. Smear / glue character. Slight LF tuck for clarity in

FACE	ATTACK	RELEASE TAIL	HF TILT	LF TILT	CHARACTER
					master-bus role.
Studio Cream	10 ms	60 ms + 3 s	−2 dB @ 5 kHz	−1 dB @ 80 Hz	Vintage warm. Slightly longer release than Silver, gentle HF + LF rolloff.

Switching face does NOT move knobs. Your Peak Reduction, Gain, Mix, Mode, and saturation settings stay where you put them. Face changes the *character* of the engine, not your sound design — switch face, hear the character difference, leave knobs alone.

For knob recall (use-case starting points like “Vocal Lead” or “Bass DI”), use the orthogonal Preset Bar — see **§11 Presets**.

> **Why per-face variance?** *Real T4B cells age. UA’s own service notes show ±30 % attack/release spread across cells of the same era. Each face here is a believable cell from the same family, not a different compressor architecture. Silver is your specific cell; the others are nearby cells you might have on a different day.*

4. COMPRESSION VS LIMIT

Per the schematic, **SW3** toggles a single 2.7 kΩ coupling resistor **R7** in or out of the path between the T4B audio-cell output and the V1A grid:

- **COMP** — R7 is **in series**. The audio sees a higher source impedance into the V1A grid, so a given photocell resistance contributes a smaller swing in attenuation. The result is a softer, more gradual onset and a lower effective ratio.
- **LIMIT** — R7 is **bypassed**. The photocell couples directly to the V1A grid. The same cell-resistance swing produces a much bigger attenuation change — the compressor engages harder and faster, behaving as a peak limiter.

The **sidechain math is identical in both modes**. PR drives the detector at the same scale either way; only the audio-side response changes. This is why flipping COMP → LIMIT at a fixed PR setting can give you noticeably more reduction without you touching the PR knob.

---

5. USE CASES

6.1 VOCALS

Start in **COMP**, **PR 30–50**, **GAIN 40**. The LA-2A is famous on vocals because the program-dependent two-time-scale photocell release gives "the voice rides forward without pumping." Watch the GR meter settle around **−3 to −6 dB** on the loud syllables. If sibilance is biting, add a touch of **EMPH** (10–25 %) to lean the sidechain into the highs — the compressor now ducks the esses on its own.

6.2 BASS

Start in **COMP**, **PR 50–70**, **GAIN 40**. Bass benefits from the slow attack — transients pass through, the body of the note rides. A small amount of **EMPH** (10–20 %) thickens the sidechain response and adds a perceived warmth on the body of the note. If the low end is triggering too much GR, dial in **SCF around 80–120 Hz** to ignore the fundamentals and let the upper-mids drive the detector.

6.3 DRUMS (PARALLEL)

Insert on a drum bus or room mics, **LIMIT** mode, **PR 70–90**, **MIX 30–50 %**. Drive it hard, then blend back to taste. The MIX fader is the dry-blend control — at 30 % you keep the natural transient attack of the drums and add a heavily-compressed parallel layer behind them for sustain and weight.

6.4 MIX BUS

COMP, **PR 20–30**, **GAIN 40**, **MIX 100 %**. Light glue across the mix. With **Stereo Link** ON the L/R channels share a detector so panned transients don't pull the image. Watch for **−1 to −2 dB** maximum GR — the LA-2A on a bus is for cohesion, not for level control.

---

6. STEREO LINK

The ∞ pill under the VU couples the L and R sidechains. **Default: OFF** — engaging stereo link is an opt-in, matching SSL/UAD/Plugin Alliance bus-comp convention. **For any stereo program material you should turn it ON** (see the table below).

HOW IT WORKS

- **Link OFF (independent)** — each channel runs its own compressor with its own sidechain. L's GR is driven by `|L|`, R's GR is driven by `|R|`. On a hard-panned loud transient (kick on left, nothing on right), L compresses heavily and R doesn't move at all.
- **Link ON (coupled)** — both channels see `max(|L|, |R|)` as their shared sidechain. Both compress by the **same** amount. The quiet channel ducks too, but since it's near silent you don't hear that part — what you hear is uniform GR with the stereo image **intact**.

WHEN TO USE WHICH

MATERIAL	LINK	WHY
Mix bus / mastering bus	**ON**	Preserves stereo image. Always.
Stereo synth pad / strings	**ON**	Same — width stable.
Stereo drum bus (default starting point)	**ON**	Smooth, no side-flicker on hard-panned transients.
Drum bus, creative "punch a hole"	**OFF**	Lets a hard-panned kick duck only its own side — tight dimensional groove. Use sparingly.
Parallel (NY) compression on stereo material	**ON**	The blend depends on consistent GR across channels.
True dual-mono (two unrelated signals on L+R)	**OFF**	They have nothing to do with each other — don't make them ride together.

WHAT YOU'LL HEAR

With link ON, the compressor sounds **smoother and wider** even though GR is often slightly **higher** than independent mode (the max-detector catches the loudest channel). The "less compressed" feel is the absence of side-flicker artifact, not lower GR.

With link OFF, the compressor sounds **more aggressive** and the stereo image **wanders** (the side with the transient ducks; the quiet side stays bright). Useful as a creative effect, wrong on a master.

INTERNALS

The link parameter is exposed as 0–100 % in APVTS so DAWs can automate between the extremes — but the front-panel ∞ button is binary (0 / 100). For surgical mid-blend (e.g. 60 % linked) automate `link` directly.

7. TIPS

1. **Start at PR 50 and walk it from there.** That's the LA-2A's working anchor — most program will land in the unit's sweet spot at PR 40–60.
2. **GAIN does not affect compression.** Set PR for the ride first, GAIN second. Treat GAIN as a clean output trim.
3. **LIMIT with low PR is your "limiter."** Don't reach for LIMIT thinking you'll only ever use it at PR 90. LIMIT at PR 30 is a gentle peak catcher; LIMIT at PR 80 is a brick.
4. **Saturation MATCH compensates for LIFT level.** When you engage LIFT, the saturator adds level (more harmonics, more energy). MATCH brings it back near the clean output so your A/B is honest.
5. **EMPH > Drive when it's an air-band sibilance issue.** A small EMPH lift (10–20 %) is more transparent than driving the compressor harder with PR.
6. **SCF on the mix bus around 60–80 Hz** stops kick + bass dominating the ride. Leave it OFF on vocals — you actually want LF to inform the body of the compression on a single voice.
7. **MIX is your "wetness" control, not your "amount" control.** PR and LIMIT control how much compression happens; MIX controls how much of it you hear.
8. **PHASE is also a sanity check.** Before committing a take, flip PHASE: if the result sounds the same, you have polarity-symmetric program. If it sounds drastically different, look for asymmetry upstream (clipping, comb filtering, polarity-flipped mics).
9. **Stereo Link on for stereo program, off for dual-mono.** If you're using one instance across L+R of a stereo bus, link it. If you're using one instance per dual-mono insert, unlink.
10. **The ∞ button moves both knobs together.** Engage it, drag GAIN, PR follows. Useful for "more compression with matching makeup" in one motion.

8. SPECIFICATIONS

ITEM	VALUE
Formats	VST3 · AU · CLAP · Standalone (macOS) — Windows x64 in active porting
Sample rate	44.1 / 48 / 88.2 / 96 / 176.4 / 192 kHz
Bit depth	32-bit floating-point internal
Channels	Mono, stereo

ITEM	VALUE
Latency	0 samples (no look-ahead, no oversampling pre-delay)
Architectures	Universal binary — Apple Silicon (arm64) + Intel (x86_64)
CPU	Approximately 1 % of one core at 96 kHz on Apple Silicon and Intel
Modeled hardware	Universal Audio LA-2A (2014 reissue)
Schematic reference	Teletronix C-10953 rev C (8 Aug 1968)
Output transformer modeled	Sowter Type 8940 (5 kΩ : 600 Ω, 3:1 step-down)
Modeled max output	+16 dBm (matching hardware spec)
Frequency response	+0 / −1 dB, 30 Hz – 15 kHz (at modeled defaults)
Calibration result	8 of 9 musical conditions (vocal/drums/bass × PR=25/50/75) match unit hardware within ±2 dB. Vocal at PR≈40–60 has a known ~−4 dB deviation under sustained material — see §5.1 Vocals for the GAIN-knob workaround. Targeted for full closure in v1.x via the Bendiksen + Najnudel 2-state ODE photocell upgrade.

9.1 DSP MODEL STATEMENT

MENTESH2A is an **analytical DSP model** — every coefficient traces to a published physics equation or a measurement on our specific unit:

- **Tube cascade stages** use Bendiksen triode waveshaper parameters (Pakarinen-Yeh-Bendiksen, DAFX §4.3.1, Eq 4.13).
- **Output stage soft-clip** via tanh with ceiling fitted to −2.18 dBFS from V2B + Sowter 8940 saturation captures (PR=0, 1 kHz, four input levels).
- **Sidechain frequency shaping** via Zölzer one-pole shelves (DAFX §5.2.4, Eq 5.31 cut variant) and RBJ biquad peak EQ (Audio EQ Cookbook).
- **Three-branch τ opto envelope follower** fitted from PR-curve measurements (R^2 = 0.85–0.97 per branch across PR = 50 / 70; PR = MAX excluded as poor-fit region).

Not a convolution sampler. Continuous behaviour across the parameter space — no interpolation between captured kernels. See `docs/POSITIONING.md § Section 4` for full per-coefficient source map.

9. TROUBLESHOOTING

The plug-in doesn't show up in my DAW. - Confirm `MENTESH2A.vst3` is in `~/Library/Audio/Plug-Ins/VST3/`. - Force a plug-in rescan (most DAWs have a "rescan" or "reset cache" button in their plug-in manager). - Quit and re-open the DAW.

macOS says the plug-in is "from an unidentified developer" or "damaged." - Right-click `MENTESH2A.vst3` in the Finder and choose **Open**. This adds a one-time gatekeeper exception. - If that doesn't help, run from Terminal: `xattr -dr com.apple.quarantine ~/Library/Audio/Plug-Ins/VST3/MENTESH2A.vst3` then rescan.

No sound. - Check **POWER** is on (bat-handle in the right cluster, label `ON`). - Check **MIX** is not at 0 % (would mean fully dry — actually you'd still hear input, so this only matters if the host is also bypassing the dry path). - Check the host channel isn't muted or routed to nowhere. - Confirm the VU is moving — if the GR meter shows activity, the audio path is alive.

A parameter isn't responding to automation. - Confirm the parameter ID matches what your DAW is writing. The MENTESH2A v1 parameter IDs are: `peakRed`, `gain`, `mode`, `bypass`, `mix`, `link`, `scHpf`, `emphasis`, `outTrim`, `phase`, `satDrive`, `satMode`, `satLow`, `satMid`, `satHigh`, `satMatch`. These IDs are stable across all v1.x versions and will never be renamed. - If you saved a session with an older build that had different IDs, the host may have orphaned automation lanes. Re-arm the lane on the current parameter and re-record.

The needle on the VU "stutters" or doesn't update smoothly. - The meter is driven by 30 Hz host events, not by audio block. If your DAW is under heavy CPU load and dropping UI events, the needle may pause briefly. The audio path is unaffected.

The plug-in sounds slightly different at 44.1 vs 96 kHz. - That is expected. The model's tube and photocell stages are sample-rate aware but, like the hardware itself, have minor resonance and rolloff differences at the band edges. The compression envelope behavior is calibrated to match within ±2 dB across all supported sample rates.

10. CREDITS

MENTESH2A is designed and built by **MENTESH ENGINEERING / Mentesh**.

The plug-in's behavior is derived from primary-source schematic and unit-capture work, not from third-party summaries. Schematic and specification sources cited:

- **S1** — Teletronix / United Recording Electronic Industries, *LA-2A Leveling Amplifier Schematic*, drawing **C-10953 rev C**, 8 August 1968, Los Angeles, California.
- **S2** — DUKA redraw of S1 (20 July 2003) with annotated tube operating voltages.
- **S3** — Universal Audio, *LA-2A User's Guide*, Rev 1.3, P/N LA2A-M01, © 2000.
- **S5** — Universal Audio product page, *Teletronix LA-2A Classic Leveling Amplifier* — modern reissue specifications.
- **S11** — E. A. Sowter Ltd / Carnhill Transformers, *Type 8940 Low- DCR Mumetal Line Output Transformer* datasheet — the T2 output transformer in the 2014 UA reissue.

Calibration was anchored to the operator's own 2014 UA LA-2A reissue via a 140-capture stimulus session (29-30 April 2026) and a fitted HF shelf at 5654.7 Hz / −6.90 dB, applied to the plug-in's voicing.

11. PRESETS

The **Preset Bar** lives in the bottom-left of the plug-in window. Two banks, both visible in the dropdown.

12.1 FACTORY BANK

Eight presets ship with the installer, covering common starting points:

PRESET	WHAT IT'S FOR
Vocal Lead	Light leveling for lead vocal — gentle PR, MIX 100 %
Vocal Aggressive	Pop / hip-hop forward vocal — heavier PR, slight EMPH
Bass DI	Direct bass — slow attack lets the pluck through
Drum Bus	Bus glue — modest PR, link ON suggested

PRESET	WHAT IT'S FOR
Kick Punch	Snappy kick / parallel kick layer
Mix Bus Glue	1–2 dB GR cohesion across the master
Parallel Smash	Heavy GR for parallel-blend layers (mix back to taste)
Acoustic Guitar	Clean leveling for picked acoustic

Factory presets install to:

```
/Library/Application Support/MENTESH ENGINEERING/MENTESH2A/Presets/Factory/
```

Each is a small JSON with the parameter snapshot. They survive plug-in upgrades (the installer overwrites the Factory dir; your User dir is untouched). To uninstall a factory preset, delete the JSON manually.

12.2 USER BANK

Click **+ Save...** in the Preset Bar to capture your current knob state. Name, description, and category. The dialog writes a JSON to:

```
~/Library/Application Support/MENTESH ENGINEERING/MENTESH2A/Presets/User/
```

User presets persist across DAW sessions, plug-in versions, and host projects. Recall is one click from the dropdown.

12.3 FACE VS. PRESET

Face and preset are orthogonal axes:

- **Face** = which cell you're on (Silver / Grey / Black / Purple / Cream). Affects *character* (cell speed, HF/LF tilt). Does **not** move knobs.
- **Preset** = which knob positions you've recalled. Affects *sound* (PR, Gain, Mix, Mode, sat). Does **not** change face.

Five faces × eight factory presets = 40 starting combinations, before you save your own.

12.4 HAND-EDITING PRESETS

Both factory and user presets are plain JSON. You can open them in any text editor and tweak values. Schema lives at `Presets/Factory/SCHEMA.md` (installed alongside the factory bank) and documents every parameter ID, range, and default. Reload the plug-in after editing to pick up changes — the Preset Bar re-scans the dirs on plug-in instantiation, not live.

12. ACTIVATION & LICENSING

MENTESH2A uses a server-verified Ed25519 license token. Three states:

13.1 TRIAL — FIRST 14 DAYS

On first instantiation the plug-in starts a **14-day full-feature trial**. The top-right banner reads **TRIAL · N DAYS REMAINING**, counting down. All five faces, both modes, all parameters — full power.

13.2 LICENSED — AFTER ACTIVATION

Click the trial / license banner to open the License panel. Paste your key (looks like `MNTSH-XXXX-XXXX-XXXX-XXXX`), click **Activate**. The plug-in connects to `me-plugins.com`, validates the key against your machine, and stores the signed token at:

```
~/Library/MENTESH ENGINEERING/MENTESH2A.lic
```

After activation the banner switches to **LICENCED · BETA** (or **LICENCED · PAID** post-launch). No more countdown.

The license is **machine-bound** — one machine per key. To move to a new machine: click **Deactivate** in the License panel first (frees the seat), then activate on the new one.

13.3 DEMO — AFTER TRIAL EXPIRES

If you don't activate within 14 days, the plug-in drops into **demo mode**: 30 seconds of full pass-through audio, 200 ms fade out, 1 second of silence, 200 ms fade in, repeat. The demo gate is per-instance and audible — it makes the plug-in usable enough to evaluate but unfit for real production work. Buy a key at https://me-plugins.com to remove.

13.4 OFFLINE / NETWORK BLIPS

Heartbeats run weekly to confirm the license is still valid. If the server is unreachable, you keep working — the plug-in falls back to **scope-aware grace**:

- Paid scope: 90 days offline grace
- Trial scope: 14 days
- Beta / comp / NFR scope: 30 days

After grace expires without a successful heartbeat, the plug-in returns to demo mode. Activate again when you're back online.

13.5 BUYING A KEY

https://me-plugins.com — check there for current pricing and bundle options. Beta testers receive a key by email after signup; production keys are issued post-launch.

13.6 PRIVACY

The activation request sends: a hash of the host machine, your license key, and a User-Agent string identifying the plug-in version. No audio, no DAW data, no project metadata. Heartbeats send the same envelope. The full handshake is documented in `source/license/LicenseManager.h`.

---

13. PLUGIN WINDOW

14.1 RESIZE

The plug-in window is **user-resizable**. Drag the bottom-right corner to scale. The 1460:720 aspect ratio is locked (you can't squash the faceplate). Limits:

- Minimum: 730 × 360 (50 % — laptop-friendly, knobs still clickable)
- Default: 1460 × 720
- Maximum: 2920 × 1440 (200 % — for 4K / 5K screens)

The React layout auto-scales based on container width, so every element stays proportional and knob hit-areas scale with the window.

14.2 PERSISTENCE

Window size is stored by the host, per plug-in instance, in the project file. Ableton Live, Logic, Reaper, Cubase, and Studio One all preserve it across re-opens. Some hosts (older Pro Tools) reset to default on each re-instance.

14.3 TWEAKS PANEL

The ⚙ **Tweaks** button at the top-right (5th tab, after DSP Architecture) toggles a side panel with visual / cosmetic settings:

- **Faceplate** — pick one of the 5 cell variants (see §3.12)
- **Knob style** — 9 cap styles (Bakelite, Studio Satin, Davies, etc.)
- **Wear** — patina / aging amount (0–100 %)
- **Tube glow** — VU window glow intensity
- **Lighting** — Bright / Dark
- **Show screws** — corner mounting hardware
- **Show saturation column** — hide the bottom saturation strip

All Tweaks settings are stored in `localStorage` in the plug-in's WebView, persistent per machine (not per project). They affect display only, not DSP.

Manual revision 2026-05-05 · MENTESH ENGINEERING · Model MENTESH2A · v1.0 Beta