



SIREN

VOCAL PITCH · THE VOICE THEY COULD NOT RESIST

USER MANUAL

Version 1.0 · VST3 + AU · macOS & Windows

Ghostnote LLC

Contents

1. Introduction	3
2. Installation & Compatibility	4
2.1 Formats and requirements	4
2.2 First install — security warnings	4
2.3 First scan	4
2.3 A recorded-takes design	4
3. Quick Start	6
3.1 Automatic correction (no editing)	6
3.2 Seeing and editing the take	6
4. How SIREN Works	7
4.1 Hearing the pitch	7
4.2 Choosing the target note	7
4.3 Retune speed done right	7
4.4 Moving the pitch without moving the voice	7
5. The Grid	8
5.1 Editing gestures	8
5.2 The status line	8
6. Controls Reference	9
6.1 Choosing a scale	9
7. Presets	10
8. Recipes	11
9. Measured Performance	11
10. Troubleshooting	12
11. Specifications	12

1. Introduction

SIREN is pitch correction built for one instrument only: the human voice. Choose the key of your song and the grid lights up the lanes that belong to it — every sung note lands on that grid, and anything off-scale is flagged red with a label saying exactly how many cents it missed by. **You can see what needs fixing at a glance, before you hear it.**



- **Two workflows in one.** Leave SIREN running as an automatic corrector (like a classic tuner), or press CAPTURE and edit the take note by note — drag a note to a new pitch, exclude one, reset another. Playback follows your edits.
- **Key-aware, not just chromatic.** Correction snaps to the scale you chose, so a slightly-flat note is pulled to the note you meant, not the nearest semitone. Seven scales are built in, plus Chromatic when you want plain semitone snapping.
- **The voice stays a voice.** Correction is formant-preserving (TD-PSOLA): pitch moves, but the character and body of the singer stay put — no chipmunk artefacts unless you ask for them (Formant Off is the deliberate effect switch).

2. Installation & Compatibility

2.1 Formats and requirements

Item	Details
Formats	VST3 (macOS and Windows) · Audio Unit (macOS)
macOS	macOS 11.0 or later — universal binary (Apple Silicon + Intel)
Windows	Windows 10 or later, 64-bit
VST3 path (macOS)	/Library/Audio/Plug-Ins/VST3/Ghostnote SIREN.vst3
VST3 path (Windows)	C:\Program Files\Common Files\VST3\Ghostnote SIREN.vst3
AU path (macOS)	/Library/Audio/Plug-Ins/Components/Ghostnote SIREN.component
Window	880 x 560 px
Latency	About 58 ms at 48 kHz (pitch lookahead), reported to the host for automatic compensation — see the note below

2.2 First install — security warnings

The installer is not yet signed with an Apple / Microsoft code-signing certificate, so both systems show a one-time warning when you first run it. The file is safe — every release is built and validated automatically on GitHub's servers — the warning only reflects the missing certificate:

- **macOS 15 (Sequoia) and later:** double-click the .pkg. When macOS says it "could not verify" the installer, close that dialog, open **System Settings > Privacy & Security**, scroll to the bottom and click **Open Anyway** next to the installer's name, then confirm.
- **macOS 14 and earlier:** Control-click (right-click) the .pkg in Finder, choose **Open**, then click **Open** in the dialog.
- **Windows:** if SmartScreen shows "Windows protected your PC", click **More info**, then **Run anyway**.

2.3 First scan

- **FL Studio (macOS & Windows):** Options > Manage plugins > Find more plugins, or press the rescan button. The plugin appears as "Ghostnote SIREN" under Installed / Effects.
- **Logic Pro (macOS):** the AU is picked up automatically on launch; if not, restart Logic or rescan in Settings > Plug-in Manager.
- **Ableton Live:** enable VST3 (and Audio Units on macOS) in Settings > Plug-Ins, then rescan.
- **Other hosts (Cubase, Studio One, REAPER...):** any VST3-capable DAW on either platform works; run the host's plugin rescan after installing.

2.3 A recorded-takes design

Accurate, natural correction needs to see a little of the future — SIREN uses roughly 58 ms of lookahead. Your DAW compensates this automatically during playback and mixdown, so recorded takes stay perfectly in time. It does mean SIREN is not meant for live monitoring while tracking: record dry, then correct. That is the intended workflow.

3. Quick Start

3.1 Automatic correction (no editing)

Step	Action
1	Insert Ghostnote SIREN on the recorded vocal track.
2	Set KEY and SCALE to your song (e.g. A Minor). If you are unsure of the key, capture the take first — the red notes will tell you when the choice is wrong.
3	Pick a preset: <i>Natural</i> for transparent correction, <i>Tight Pop</i> for modern pop, <i>Hard-Tune</i> for the robotic snap, <i>Rap Melodic</i> for tuned rap.
4	Play. Correction is live from the first note.

3.2 Seeing and editing the take

Step	Action
1	Press CAPTURE , play the vocal once from the top, press STOP .
2	Every sung note is now a block on the grid. Ivory = fine; red = off-scale, with a cents label.
3	Fix what you disagree with: drag a note up/down to retarget it (it turns teal), click a note to exclude it from correction, double-click to reset it to automatic.
4	Play back — correction follows the grid, including your edits, anchored to the song timeline.

4. How SIREN Works



the correction is smoothed, not the target — note changes stay instant, no portamento

4.1 Hearing the pitch

SIREN tracks the voice with the YIN algorithm — the standard for monophonic vocal pitch — analysing every 5 ms across the vocal range (60 Hz to 1.1 kHz). On steady notes the measurement is accurate to better than 3 cents (verified in the release test suite). Unvoiced sounds (breaths, consonants, silence) are detected as unpitched and passed through untouched.

4.2 Choosing the target note

Each detected pitch is snapped to the nearest note of your KEY and SCALE. A note you retargeted on the grid overrides the automatic choice while the playhead is inside it, and an excluded note is left completely alone. Chromatic scale turns key-awareness off and snaps to the nearest semitone — useful for melodies that deliberately leave the key.

4.3 Retune speed done right

Classic correctors smooth the *target*, which turns every interval into a small portamento — the giveaway "tuner slide". SIREN smooths the *correction* instead: when you sing a clean note change, the new note sounds immediately; only the drift and scoop *within* a note is flattened at the RETUNE speed. The release tests verify a clean two-semitone jump arrives on the new pitch immediately even at slow RETUNE settings.

4.4 Moving the pitch without moving the voice

Correction is applied with TD-PSOLA, the classic technique for natural vocal retuning: the voice is rebuilt from pitch-synchronous grains, so the spectral envelope — the formants that make the singer sound like themselves — stays in place while the pitch moves. Setting FORMANT to Off resamples the grains instead, shifting character along with pitch: the deliberate chipmunk/deep-voice effect.

5. The Grid

The grid is a piano roll of your captured take. Lanes are pitches; blocks are the notes you sang; the gold vertical line is the playhead.

	On the grid	Meaning
	Gold lanes	Notes that belong to the chosen key and scale.
	Dark lanes	Out-of-key notes.
	Ivory block	A sung note that sits on scale — left alone.
	Red block	Off-scale (or more than 35 cents out); the label shows how far, e.g. "+37c".
	Teal block	Manually retargeted (you dragged it); a faint ghost marks the original pitch.
	Struck-through block	Excluded — SIREN will not correct this note.

5.1 Editing gestures

- **Drag a block up or down** — retarget the note to any lane, in or out of the scale. It turns teal, and a ghost outline remains at the sung pitch.
- **Click a block** — toggle exclusion. Excluded notes are struck through and never corrected (keep a stylistic slide or a spoken word untouched).
- **Double-click a block** — reset it to fully automatic.

5.2 The status line

Under the controls SIREN reports what it knows: while capturing, "listening... play the vocal from the top"; after a capture, the note count and how many are off-scale; when idle with no capture, a reminder that correction is live anyway. The count of red notes is a fast key-check: if half the take reads off-scale, the KEY or SCALE is probably wrong, not the singer.

6. Controls Reference

Control	Range	Default	What it does
KEY	C – B	A	The root of your song. The grid lights the lanes that belong to this key.
SCALE	7 scales	Minor	Which notes count as in key: Major, Minor, Dorian, Mixolydian, Harmonic Minor, Melodic Minor, Chromatic (= snap to any semitone).
RETUNE SPEED	0–100 %	25 %	How fast pitch is pulled to the note. 0 = instant robotic hard-tune; higher = slower, more natural correction that preserves vibrato.
AMOUNT	0–100 %	100 %	How far off-pitch notes are pulled toward the target. 100 % = all the way.
FORMANT	Off / Preserve	Preserve	Preserve keeps the natural character of the voice while pitch moves. Off shifts character with pitch (the deliberate effect).
MIX	0–100 %	100 %	Blend of corrected and original voice.
CAPTURE	—	—	Press, play the take once from the top, press STOP — the take appears on the grid for editing.
PRESET	4 programs	—	Correction characters (sets RETUNE / AMOUNT / FORMANT / MIX only — never your key, scale or grid edits).
BYPASS	on/off	off	Hear the untouched voice, latency-matched and click-free.

6.1 Choosing a scale

Minor and Major cover most songs. Dorian and Mixolydian suit soul, funk and folk melodies that keep landing on the "wrong" sixth or seventh. Harmonic and Melodic Minor catch the raised sevenths of dramatic minor hooks. When a melody deliberately borrows outside notes, either switch to Chromatic or keep your scale and rescue the borrowed notes individually on the grid (drag or exclude) — that is usually the better-sounding fix.

7. Presets

Programs set the correction character only; KEY, SCALE and your grid edits are never touched.

Preset	Retune	Amount	Character
Natural	65 %	85 %	Transparent clean-up — drift is fixed, vibrato and phrasing survive. Nobody hears a tuner.
Tight Pop	25 %	100 %	Modern pop-tight. Clearly produced, still human.
Hard-Tune	0 %	100 %	Instant snap — the robotic hard-tune effect.
Rap Melodic	8 %	100 %	Fast snap tuned for melodic rap deliveries — hard enough to stylise, quick enough for dense phrasing.

8. Recipes

- **Transparent lead vocal:** *Natural*, correct key and scale, FORMANT Preserve. Then capture the take and exclude any expressive slides you want to keep — the tuner disappears entirely.
- **Modern pop lead:** *Tight Pop*. If a big interval jump ever sounds too eager, raise RETUNE toward 35–45 % — note changes stay instant by design, so RETUNE only shapes behaviour within notes.
- **Trap / hyperpop hard-tune:** *Hard-Tune*, and pick the scale deliberately — the effect exposes the grid, so Harmonic Minor vs Minor is an artistic choice you will hear.
- **Tuned rap:** *Rap Melodic*. Consider AMOUNT 80–90 % so pitchy talk-singing keeps some spoken texture.
- **Backing stacks:** harder settings than the lead (*Tight Pop* or RETUNE 10–20 %) lock the stack together; small per-voice errors multiply in a stack.
- **Creative reharmonisation:** capture, then drag notes to new targets — SIREN will re-sing the take. Two instances on duplicated tracks with different drags = instant harmony line, with formants intact.
- **In the Ghostnote chain:** **SIREN** first (fix pitch before anything colors or compresses the voice) → TALISMAN → VEIL → MONOLITH.

9. Measured Performance

Every SIREN build must pass an automated audio-measurement suite before release. Shipping v1.0 results:

Test	Result
Pitch detection accuracy (steady tones, 110–880 Hz)	Within 3 cents; vibrato tracked; noise and silence correctly read as unpitched
Correction (note sung 30 cents flat, hard-tune)	Lands within 5 cents of the target note, with level preserved within 2 dB
Already-in-tune input	Passes essentially unchanged (within 5 cents)
Note-change behaviour (clean jump at slow RETUNE)	Arrives on the new note immediately — no portamento smear across the interval
Formant Off mode	Still lands within 5 cents of the target
Capture segmentation (3-note melody)	3 of 3 notes found, pitch medians within 0.3 semitone; retarget / exclude / reset verified
Bypass	Bit-exact (input delayed by exactly the reported latency)
Host validation	Apple auval (strict) passed; pluginval strictness 10 passed on Apple Silicon and Intel

10. Troubleshooting

Symptom	Cause & fix
Sounds robotic when it should not	RETUNE too fast for the material. Raise it (or load <i>Natural</i>). Remember 0 % is the intentional hard-tune effect.
Correction pulls to wrong notes	Wrong KEY or SCALE — capture the take and look at the red count. Fix the key, or switch to Chromatic, or drag the disputed notes on the grid.
Vibrato is being flattened	Raise RETUNE (slower) and/or lower AMOUNT to ~85 %. Wide, slow vibrato reads as drift at fast settings.
Chipmunk / character change	FORMANT is Off. Set it to Preserve (the default) for natural correction.
Grid is empty	The grid only shows a captured take: press CAPTURE, play the vocal once from the top, press STOP. Correction itself works with or without a capture.
Edits do not seem to apply	Grid edits are anchored to the song timeline — make sure you are playing the same region of the song that was captured, from the DAW timeline (not a clip preview).
Feels late when singing through it	SIREN uses ~58 ms of lookahead by design — it is a recorded-takes corrector. Record dry and monitor without SIREN; playback is compensated automatically.

11. Specifications

Specification	Value
Detection	YIN (CMNDF) with parabolic refinement, 60–1100 Hz, ~5 ms hop
Correction	TD-PSOLA, formant-preserving (or grain-resampling with Formant Off); shift range half to double pitch
Retune law	Correction-domain smoothing, ~2 ms (0 %) to ~500 ms (100 %); note changes pass instantly
Scales	Major, Minor, Dorian, Mixolydian, Harmonic Minor, Melodic Minor, Chromatic — all 12 keys
Capture	Host-timeline anchored; up to 4000 notes; 60 ms minimum note
Stereo behaviour	One shared grain schedule for both channels — stereo-coherent correction
Latency	~2804 samples at 48 kHz (~58 ms), reported to host
Sample rates	All standard rates (analysis windows recomputed per rate). Pitch detection cost grows with rate: 44.1/48 kHz recommended for realtime playback; 88.2/96 kHz workable on modern CPUs; 176.4/192 kHz best rendered offline

Specification	Value
Formats	VST3 (macOS & Windows), AU (macOS)
Plugin codes	Manufacturer Ghst · Plugin Sirn (stable across all versions)

Ghostnote SIREN v1.0 — © Ghostnote LLC. Support: support@ghostnotesupply.com. VST is a trademark of Steinberg Media Technologies GmbH. Audio Units is a trademark of Apple Inc.