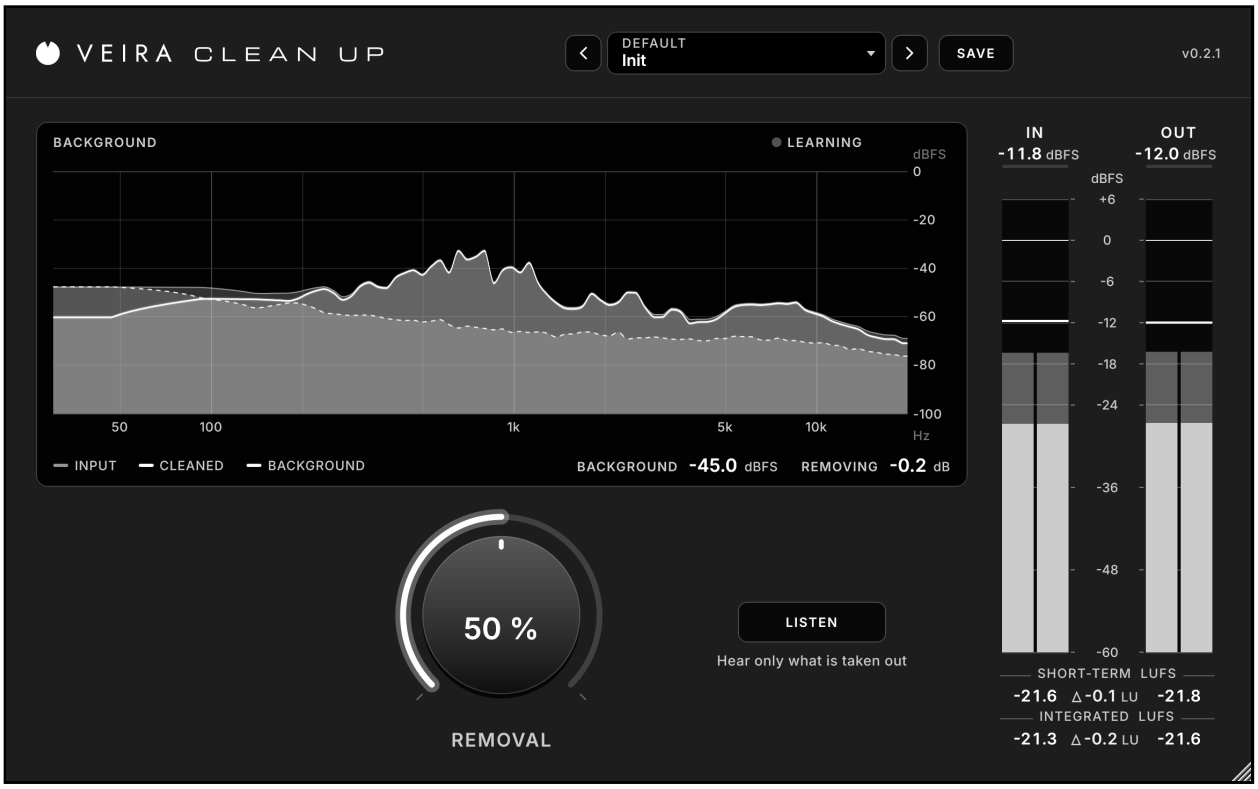
VEIRA

BACKGROUND SOUND REMOVER

VEIRA Clean Up

User Manual

Take the room tone, hiss, hum, fans and traffic out from under a voice or an instrument, with one knob. The plug-in learns the background by itself, in the pauses, and LISTEN lets you hear exactly what goes.



Contents

Welcome 3
Getting started 4
What you need 4
Install with VEIRA Manager 4
Activation and the licence 4
Formats and where the files go 5
The panel at a glance 6
Using the controls 7
How background removal works 8
The background is what stays when the voice stops 8
Where the voice is, and where it isn't 8
How it tells a pause from a word 9
What happens inside VEIRA Clean Up 11
Latency, bypass and bouncing 12
Putting it together: a clean voice in four moves 12
Controls 13
REMOVAL and LISTEN 13
Displays 15
The BACKGROUND display 15
Meter bridge (IN | OUT) 17
Presets 19
The preset browser 19
The preset menu 19
Saving your own presets 19
Presets in your DAW 20
The factory folders 20
Recipes 22
A podcast episode 22
A voice-over that must sit on silence 22
A vocal recorded at home 22
An amp that hums 22
An interview recorded outside 23
A sample or an old tape transfer 23
Specifications 24
Troubleshooting and questions 25

Welcome

Thank you for choosing **VEIRA Clean Up**. It takes the **background** out of a recording: the room tone, the hiss of a preamp or a tape, the hum of a mains supply or an amp, a computer fan, air-conditioning, the traffic outside the window. What stays is the sound you recorded, the voice or the instrument, at its own level and with its own tone.

There is one knob. REMOVAL sets how much of the background goes: a little, for a good recording that only needs a polish, or all of it, for a voice-over that must sit on silence. You never have to show the plug-in what the background sounds like. It **learns it by itself**, in the pauses, and keeps that picture up to date while the recording plays. LISTEN plays only what is being taken out, so you can check that no part of the voice is going with it.

VEIRA Clean Up is made for podcasters and producers: for a podcast or a voice-over recorded at home, a vocal with a fan in the room, a guitar amp that hums between the notes, an interview recorded outside, a sample with noise before and after it, or an old transfer from tape.

What makes it different

- **One knob.** REMOVAL is the whole adjustment: how far the background goes down. Everything else, finding the background and keeping the voice out of the way, the plug-in does itself.
- **It learns in the pauses.** The background is whatever is left when the voice stops. The plug-in listens for those moments, updates its picture of the background there, and never learns from the voice itself.
- **It leaves sound without pauses alone.** A held note, a strummed guitar or a dense mix has no background to learn from, so it passes through unchanged. The plug-in only ever works where it has heard a pause.
- **Nothing is touched at 0 %.** With REMOVAL at 0 % the output is the input, bit for bit, only delayed by 21 ms. Every bit of cleaning is in the difference you ask for.
- **You hear and see what goes.** LISTEN plays only the background that is being removed. The **BACKGROUND** display draws the input, the cleaned output and the background the plug-in has learnt, with its level and how much is being taken out right now.
- **27 presets** sorted by what is being cleaned: speech, vocals, instruments, location recordings, music, and restoration jobs.
- **A legible, scalable window.** Large text, a display with scales and read-outs you can read at a glance, and a window you can scale from 75 % to 200 %.

How it differs from VEIRA Vocal Remover and VEIRA Gate. VEIRA Vocal Remover takes the *voice* out of a finished song, to make a karaoke track or an a cappella. VEIRA Clean Up does the opposite job: it takes the *background* out from under a voice. A gate, such as VEIRA Gate, mutes the track between the words and leaves it alone while the voice plays, so the background is still there under every word. VEIRA Clean Up cleans under the words too, and it never cuts the voice off: breaths and the tails of words stay.

How this manual is organised. *Getting started* covers installing and activating. *The panel at a glance* names every part of the window. *How background removal works* explains, without maths, what the plug-in listens for, how it tells a pause from a word, and why it needs a pause before it can start. *Controls* explains REMOVAL and LISTEN with drawings, *Displays* teaches you to read the BACKGROUND display and the meter bridge, and *Presets*, *Recipes*, *Specifications* and *Troubleshooting* follow.

Getting started

What you need

macOS	macOS 11 Big Sur or later. Native on Apple silicon and Intel Macs (universal).
Windows	Windows 10 or 11, 64-bit.
Plugin formats	VST3 (macOS and Windows), Audio Units (macOS), and a standalone app. Pro Tools (AAX) versions arrive once they are signed for Pro Tools.
Hosts	Any DAW that loads VST3 or Audio Units: Logic Pro, Cubase, Nuendo, Ableton Live, Studio One, Reaper, FL Studio, Bitwig and others.
Account	A VEIRA account at veiraaudio.com, and the free **VEIRA Manager** app.
Internet	Needed to install and to refresh the licence. The plugin itself never goes online.

Install with VEIRA Manager

VEIRA plugins are not downloaded from the website one by one. You install them with **VEIRA Manager**, the free VEIRA app that installs, updates and activates everything you own. Its own manual covers it in full; the short version is:

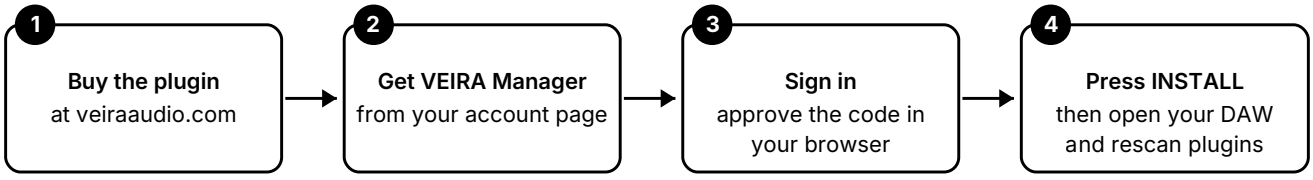


FIGURE 1 From purchase to first sound in four steps.

1. **Buy** VEIRA Clean Up at **veiraaudio.com**. It's a one-time purchase: you own it for good, with free updates.
2. **Download VEIRA Manager** from your account page (or veiraaudio.com/manager) and install it.
3. **Sign in.** VEIRA Manager shows a short code and opens your browser; confirm the code there. Your password is never typed into the app.
4. Find VEIRA Clean Up in the list and press INSTALL. When it says SHOW, it's installed. Open your DAW and let it rescan its plugins.

After installing, VEIRA Clean Up appears in your DAW under the manufacturer **VEIRA** (in Logic Pro: Audio Units ▸ VEIRA ▸ VEIRA Clean Up).

Activation and the licence

VEIRA Manager keeps a licence on your computer that the plugins check. It is refreshed every time VEIRA Manager opens and stays valid for 30 days offline, so you can work without internet as long as you open VEIRA Manager now and then.

Without a valid licence, VEIRA Clean Up still loads and works, but **its sound fades out for 1.5 seconds every 30 seconds**, and a strip along the bottom of its window says why (for example, *NOT ACTIVATED ON THIS COMPUTER. SIGN IN WITH VEIRA MANAGER.*). Clicking the strip opens your account page.

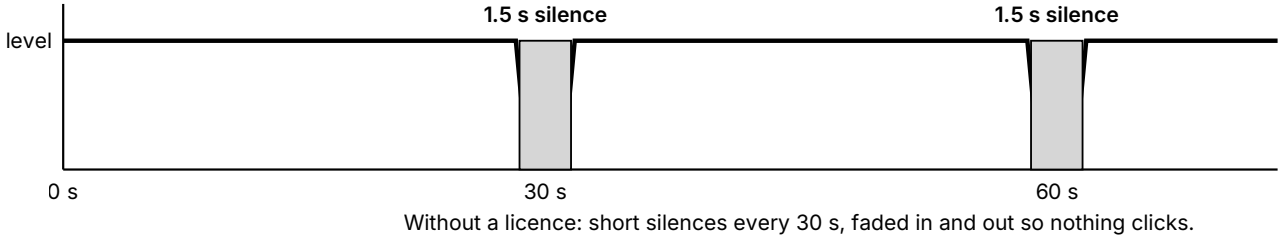


FIGURE 2 What an unlicensed copy does. With a valid licence the audio is never touched.

- **Activating while your DAW is open** takes effect within about two minutes. You don't need to reload the plugin.
- **"Not activated on this computer"**: open VEIRA Manager and sign in.
- **"Expired"**: open VEIRA Manager while online; it renews the licence at once.
- **"Not included"**: the signed-in account doesn't own VEIRA Clean Up. Check that you're signed in with the account you bought it with.

Formats and where the files go

VEIRA Manager puts every format in the standard place, so you never need to copy files by hand:

Format	macOS	Windows
VST3	`~/Library/Audio/Plug-Ins/VST3` (just for me, the default) or `/Library/Audio/Plug-Ins/VST3` (all users)	`C:\Program Files\Common Files\VST3` (all users, the default) or `%LOCALAPPDATA%\Programs\Common\VST3`
Audio Units	`~/Library/Audio/Plug-Ins/Components` or `/Library/Audio/Plug-Ins/Components`	(Audio Units are macOS only)
Standalone	`~/Applications/VEIRA` or `/Applications/VEIRA`	`C:\Program Files\VEIRA` or `%LOCALAPPDATA%\Programs\VEIRA`

You choose *just for me* or *all users* in VEIRA Manager's install settings.

The standalone app runs VEIRA Clean Up on its own, with your audio interface's input and output, which is handy for trying presets without a DAW.

The panel at a glance

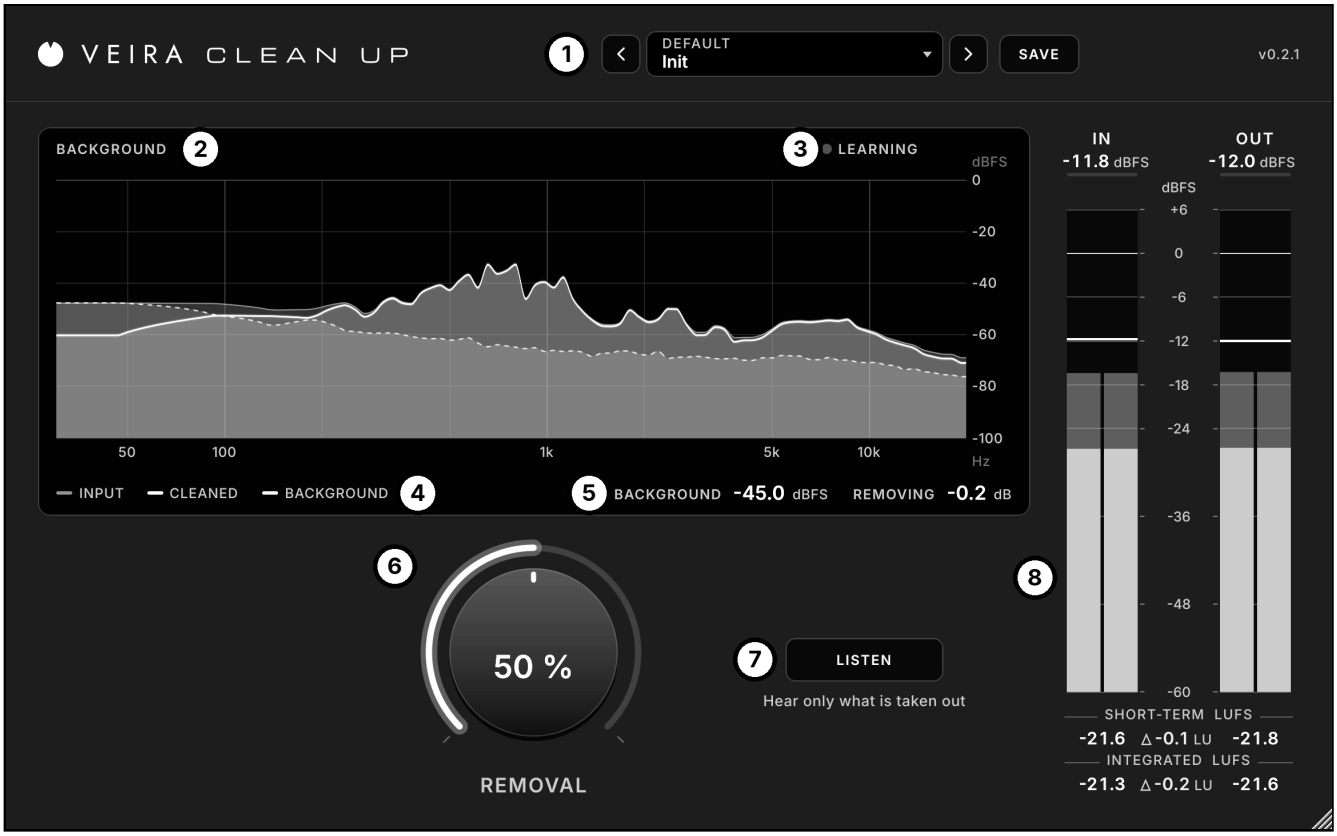


FIGURE 3 VEIRA Clean Up with the factory preset *Default: Init* loaded (REMOVAL 50 %) while a voice plays over hiss and hum. In the display the voice shows as the peaks in the middle; the background the plug-in has learnt is the dashed line underneath.

	Part	What it does
1	Preset browser	Steps through and saves presets. See *Presets*.
2	BACKGROUND display	The input, the cleaned output and the background the plug-in has learnt, 30 Hz to 20 kHz.
3	LEARNING lamp	Lights while the plug-in hears a pause and updates its picture of the background.
4	Legend	Names the three traces: INPUT, CLEANED, BACKGROUND.
5	Read-outs	BACKGROUND, the background's level in dBFS, and REMOVING, how much quieter the output is than the input right now.
6	REMOVAL	How much of the background is taken out.
7	LISTEN	Hear only what is being taken out.
8	Meter bridge	Input and output level and loudness.

The window has three areas: the **BACKGROUND display** across the top, REMOVAL and LISTEN under it, and the **meter bridge** down the right.

Using the controls

Every VEIRA plugin is played like a piece of hardware. The same few parts appear on every panel:

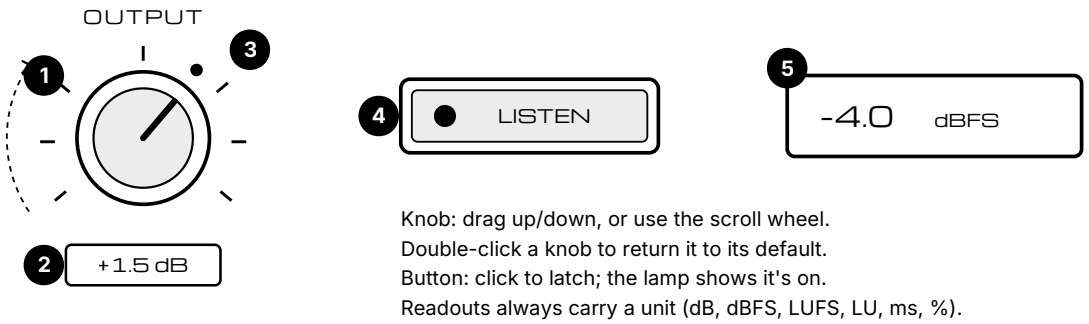


FIGURE 4 The parts you'll find on every VEIRA panel (the labels are examples).

- 1 **Knob.** Drag up to turn it clockwise, down to turn it back, or use your mouse wheel. **Double-click** returns it to its default. The engraved pointer and the small dot on the scale show where it's set.
- 2 **Value window.** The backlit window under each knob shows its exact value with its unit.
- 3 **Silkscreen label.** The control's name, as it appears in your DAW's automation list.
- 4 **Push button.** Click to switch on; the lamp lights and the cap sits lower in its bezel. Click again to switch off.
- 5 **Readout.** Live numbers update ten times a second so you can read them while the meters move. A value over 0 dBFS shows as a warning.

Tooltips: hover over any control or display for a short explanation. On displays, the tooltip depends on where you point.

Automation: every knob and button is a host parameter, so you can automate it in your DAW and control it from a MIDI controller or control surface that your DAW maps.

Units: VEIRA never shows a bare number. **dBFS** is level against digital full scale (0 dBFS is the loudest a digital signal can be); **dB** is a change in gain; **LUFS** is loudness as streaming services and broadcasters measure it; **LU** is a difference in loudness (1 LU = 1 dB).

On VEIRA Clean Up's panel the parts look a little different from the drawing above:

- **The REMOVAL knob** shows its value in its middle and its name under it; a lit ring (blue on screen; in print, the lighter arc) shows how far it is turned. Double-click (or Alt/Option-click) for 50 %.
- **Type a value:** right-click the knob (Control-click on a Mac), type a number such as **35** and press Return. Escape leaves it unchanged.
- LISTEN has no lamp: when it is on, the whole button fills (orange on screen; in print, solid with dark lettering).
- **Resize the window** by dragging its bottom-right corner: from 75 % to 200 %, proportions kept, the size saved with the session.

How background removal works

If you have used noise reduction before, some of this will be familiar; if not, these pages explain what the plug-in can and can't do, and why it behaves as it does.

The background is what stays when the voice stops

Every recording has two parts: the sound you wanted, and the **background** that was in the room with it. The background has a useful property: it is **steady**. A hiss, a hum, a fan or the traffic outside sound much the same from one second to the next, while a voice changes all the time, from word to word and from vowel to consonant. And a voice **pauses**. Between phrases, between sentences, before the first word, there are moments when the microphone hears only the room.

VEIRA Clean Up uses both facts. It listens for the pauses, and in each one it takes a picture of the background: how loud it is at every frequency, from the hum at the bottom to the hiss at the top. Then, while the voice plays, it compares the sound at each frequency with that picture. Where the sound is well above the picture, that's the voice, and it is kept. Where it sits at the picture's level, that's background, and it is turned down by as much as REMOVAL asks.

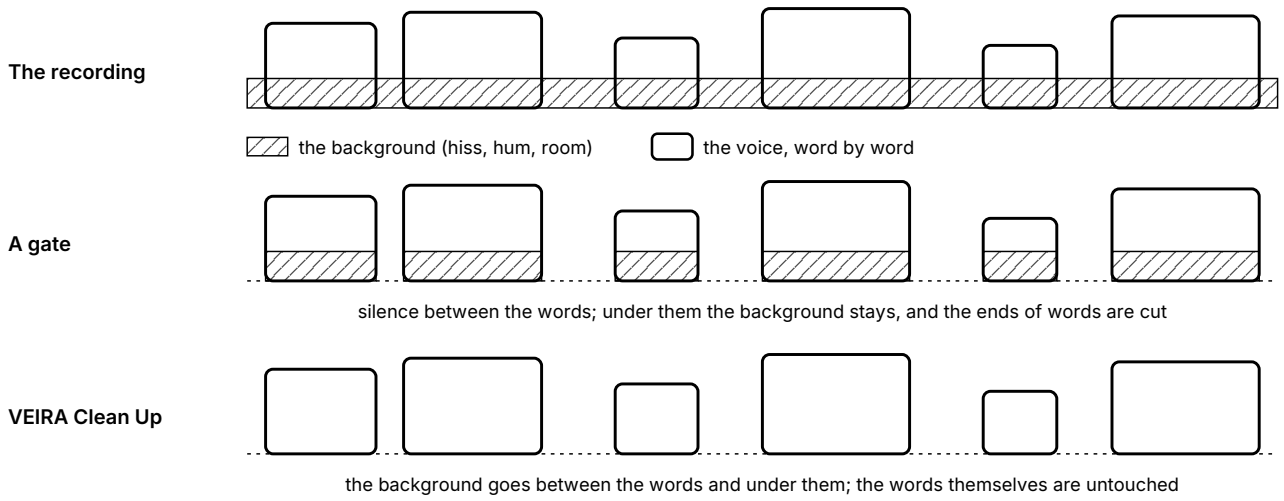


FIGURE 5 A recording, what a gate does to it and what VEIRA Clean Up does. Time runs left to right; the hatched strip is the background, the outlined blocks the words.

That is also the difference from a **gate**. A gate works on level alone: it opens for the words and closes between them, so the background is still there under every word and the quiet ends of words get cut off. VEIRA Clean Up works frequency by frequency, so it takes the background out from under the words too, and it never closes on the voice: a breath, a soft consonant or the tail of a word stays because it sits above the picture of the background.

Where the voice is, and where it isn't

The drawing below shows one moment of a spoken voice over hiss and hum, as a spectrum: frequency from left to right, level from bottom to top.

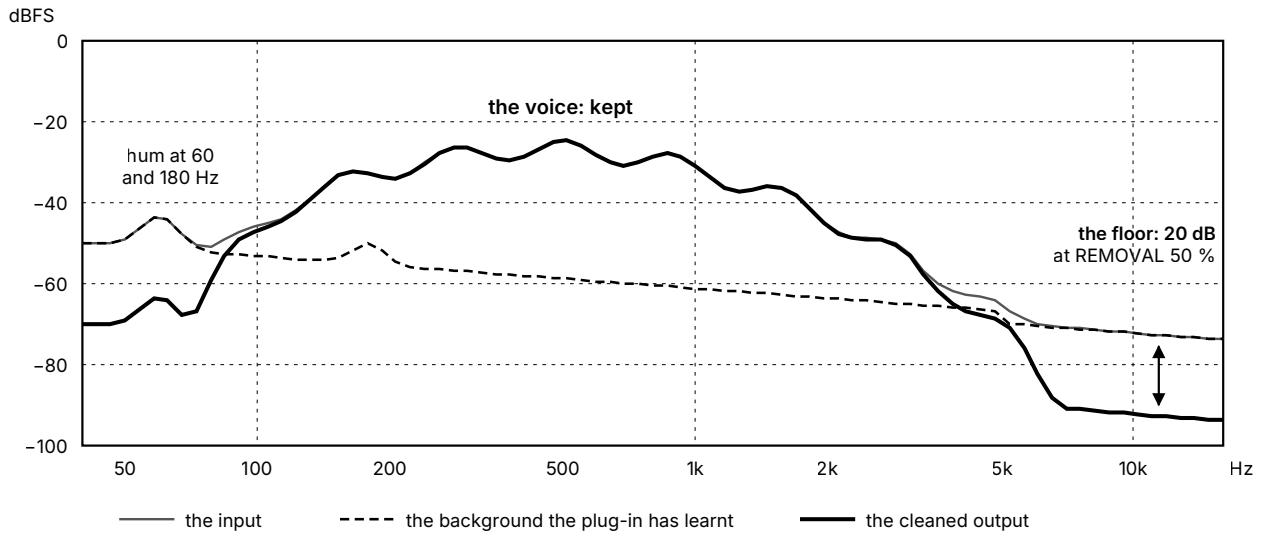


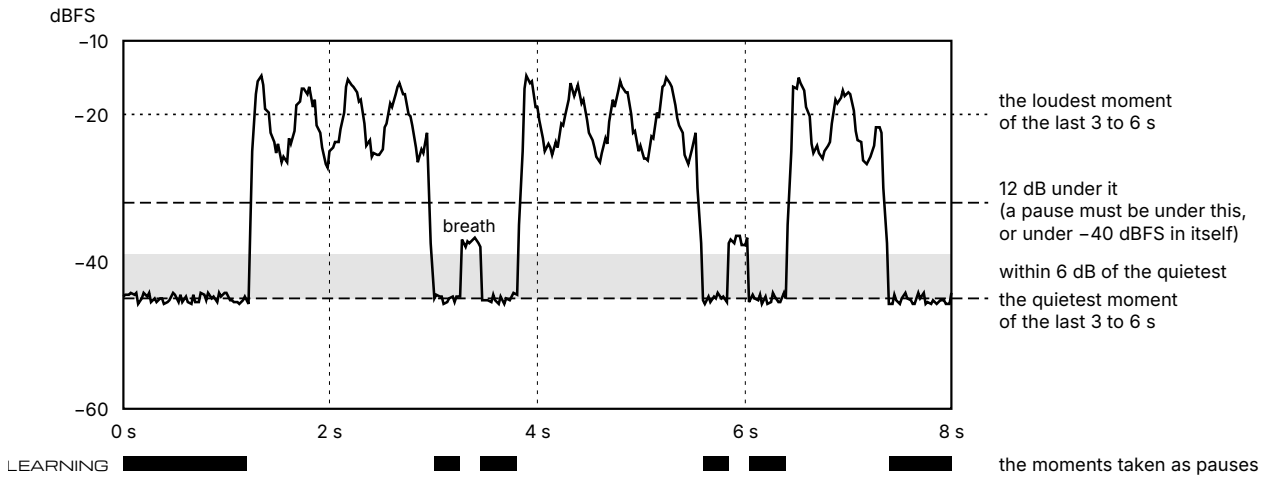
FIGURE 6 One moment of a voice over hiss and hum, at REMOVAL 50 %. Where the input rises well above the learnt background, the output follows the input; where the input is only background, the output sits 20 dB under it.

- **The input** (the thin grey line) is the voice and the background together. The voice is the broad hump in the middle, between about 150 Hz and 4 kHz; the two bumps at the left are the mains hum and its first harmonic; the slope that falls away towards the right is hiss.
- **The learnt background** (dashed) is the plug-in's picture of the room, taken in the last pause. It is a line, not a number: the hum is loud and the hiss quiet, and the picture knows both.
- **The cleaned output** (the heavy line) follows the input wherever the voice is clearly above the background, and drops to the **floor** wherever there is only background. The floor is set by REMOVAL: 20 dB under the background at 50 %.

In between, where the voice is only a little above the background (the edges of the hump, the quietest consonants), the plug-in takes out part of the background and keeps the rest, rather than deciding sharply. That is what keeps the voice from sounding chopped or watery at sensible settings.

How it tells a pause from a word

The whole method stands or falls on finding the pauses, so the plug-in is careful about them. It measures the level of the sound about 190 times a second, in the range where voices and instruments live (100 Hz to 8 kHz), and keeps track of the quietest and the loudest moments of the last three to six seconds. A moment counts as a **pause** when it is within **6 dB of the quietest** moment of that stretch, *and* either **12 dB under the loudest** or **quiet in itself**, under −40 dBFS.



The black bars mark the pauses: the lead-in before the first phrase, the gaps between phrases, the tail.
The breath in each gap rises out of the shaded band, so it is not learnt as background.

FIGURE 7 The level of a spoken voice over eight seconds, and the moments the plug-in takes as pauses. The LEARNING lamp on the display lights at exactly those moments.

A few things follow from this rule, and they are worth knowing:

- **A voice is never learnt as background**, because while it plays it is far above the quietest moment. Even a breath or the tail of a word in a gap rises out of the 6 dB band and is left out of the picture. Within a pause the plug-in is careful bin by bin too: any frequency that sits well over its own recent minimum is treated as sound, not background.
- **Sound without pauses is left alone.** A held organ note, a strummed guitar that never stops, a dense mix: nothing in them ever drops 12 dB under the loudest moment, so no moment counts as a pause, the plug-in never learns a background, and the sound passes through unchanged, at any REMOVAL. This is by design: a plug-in that guessed a background from sound that never stops would eat the sound.
- **The picture is kept between pauses.** While the voice plays, the plug-in applies the last picture it took and waits for the next pause to refresh it. If the background changes slowly (a fan speeding up, traffic building), the picture follows it pause by pause, within a tenth of a second of each.
- **Quiet backgrounds are learnt at once; loud ones at the first pause after the voice begins.** A background under −40 dBFS counts as a pause in itself, so a recording that starts with a second of room tone is learnt from its first moment. A louder background (an interview in the street at −35 dBFS, for example) is only recognised as a pause once there has been something 12 dB louder to compare it with: the first words. Until the plug-in has heard a pause it knows no background and changes nothing; the BACKGROUND read-out shows --.

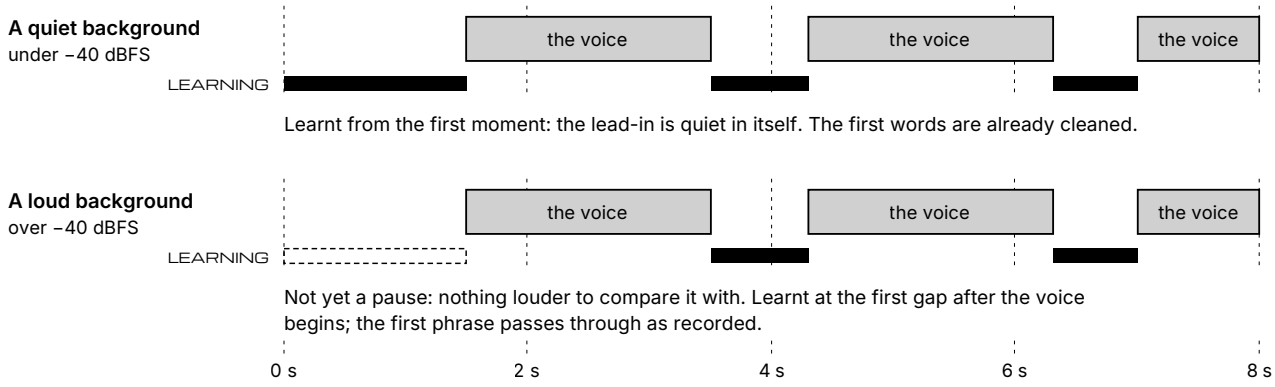


FIGURE 8 When the plug-in first knows the background. With a loud background, give it a pause early: play from a little before the first line, or leave a gap after the first sentence.

What happens inside VEIRA Clean Up

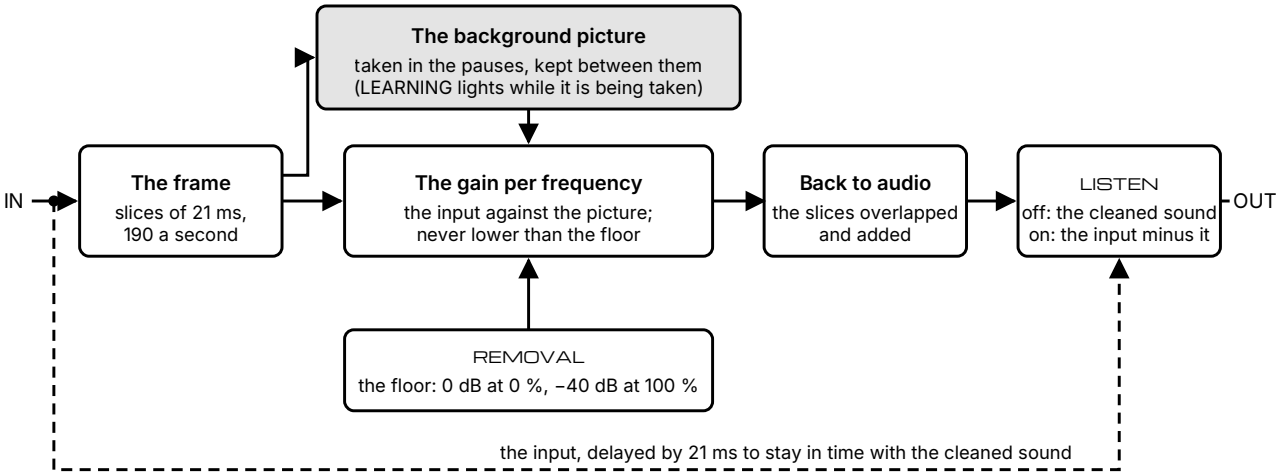


FIGURE 9 Signal flow. The sound is cut into short overlapping slices; each is compared, frequency by frequency, with the picture of the background taken in the pauses, turned down where it is only background, and put back together.

Step by step:

1. **The frame.** The sound is cut into slices of 1,024 samples (21 ms at 48 kHz), a new one every 256 samples, overlapping. Each slice is turned into a spectrum: how loud the sound is at each of about 500 frequencies. A stereo signal is analysed as one, so both channels always get the same treatment and the stereo image never moves.
2. **The background picture.** When the slice counts as a pause (see above), the picture of the background follows it, frequency by frequency, within a tenth of a second. When it doesn't, the picture is kept as it is.
3. **The gain per frequency.** Each frequency of the slice is compared with the picture. Where it is clearly above it, the gain is close to 0 dB; where it sits at the picture's level, the gain drops to the **floor** that REMOVAL sets. In between, the plug-in estimates how much of what it hears is voice and keeps that share, smoothly from one slice to the next and from one frequency to its neighbours, so nothing flickers.

Where it hears no voice at all the gain is *exactly* the floor, which is why the background never breaks up into the watery, bubbling residue (“musical noise”) that older noise reducers leave.

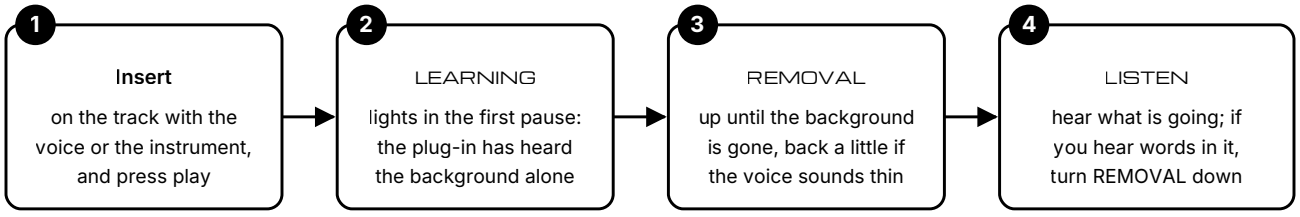
4. **Back to audio.** The slices are put back together with overlapping windows, which reconstruct the sound exactly: at REMOVAL 0 % the output is the input, bit for bit, only delayed by one slice.
5. LISTEN. Off, you hear the cleaned sound. On, you hear the input minus the cleaned sound: exactly what was taken out.

Latency, bypass and bouncing

The one slice of delay is the plug-in’s **latency**: 1,024 samples at 44.1 and 48 kHz (21 ms at 48 kHz, 23 ms at 44.1 kHz), 2,048 at 88.2 and 96 kHz, 4,096 at 176.4 and 192 kHz. It is constant and the plug-in reports it to your DAW, which **compensates** for it on playback and in a bounce, so the track stays in time with every other track. You don’t need to do anything.

- **Bypass stays in time.** When you bypass the plug-in, it still delays the sound by the same amount, so switching bypass on and off never makes the track jump. While bypassed it keeps listening, so when you switch it back in it still knows the background.
- **Recording and live monitoring.** Compensation can’t make a live sound arrive early: if a singer or a speaker listens to themselves *through* VEIRA Clean Up while recording, they hear themselves 21 ms late, which many people notice. Monitor the microphone directly, or through your interface’s own monitoring, and use the plug-in on playback.
- **Bounces** (export, bounce in place, render, freeze) are the same as playback: the plug-in learns from the material as it goes, so a bounce that starts at the beginning of a clip cleans it exactly as you heard it. LISTEN switches itself off when the host renders offline, so it can never end up in an export by accident; in a real-time bounce it stays as you left it.

Putting it together: a clean voice in four moves



A pause of about a second is all it needs to start. Most voices sit well between 30 and 60 %.

FIGURE 10 The quickest way to a good result. Or load a preset from the folder that matches your recording: each one is a REMOVAL setting for a situation.

Controls

REMOVAL and LISTEN

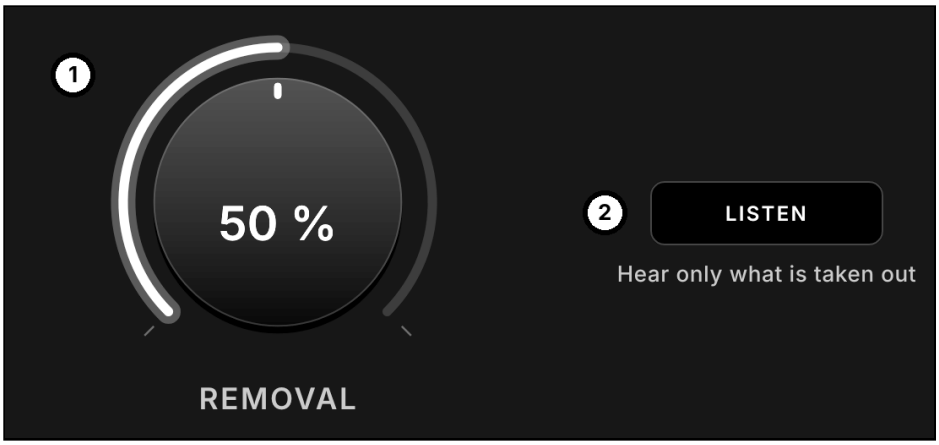


FIGURE 11 1 REMOVAL, with its value in its middle and the lit ring showing how far it is turned, and 2 LISTEN, off here, with its reminder line underneath.

Control	Range	Default	What it does
REMOVAL	0 to 100 %, in 1 % steps	50 %	How far the background is taken down: 10 dB at 25 %, 20 dB at 50 %, 40 dB at 100 %. 0 % leaves the sound exactly as it is. Right-click to type a value.
LISTEN	On, off	Off	Hear only what is being taken out. Not saved in presets, not automatable, off when the host renders offline.

REMOVAL

REMOVAL sets the **floor**: the most the plug-in will turn any frequency down. The floor is a straight line, 0.4 dB for every percent: 0 dB at 0 %, −10 dB at 25 %, −20 dB at 50 %, −30 dB at 75 % and −40 dB at 100 %.

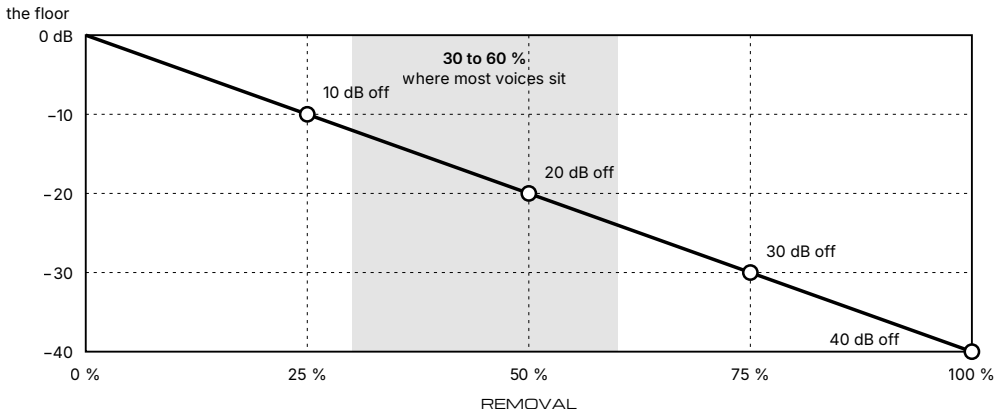


FIGURE 12 The floor against REMOVAL. It is what you hear in the pauses; under the voice the plug-in takes less, because part of what it hears there is voice.

What you hear. In the pauses, the background goes down by exactly the floor: at 50 % a hiss at −45 dBFS becomes a hiss at −65 dBFS. Under the voice, only what is clearly background goes, so the voice keeps its level and its tone: on a spoken voice over a background 33 dB quieter, the voice's level moves by less than half a decibel at 50 %. Breaths and the tails of words are part of the voice and stay, so the gaps between a speaker's words end up a little above the floor rather than on it.

Where to set it. Most voices sit well between **30 and 60 %**: enough to take a home studio's fan, computer and room out of the way (12 to 24 dB) without touching the voice. Turn it up until the background is gone, then back a little if the voice starts to sound thin or watery. Three rules of thumb:

- **The quieter the background, the more you can take.** A good recording with a little air can go to 100 % and sound like silence between the lines, because there is nothing left to hear. A loud background close under the voice (10 dB under it, as from a camera's microphone outside) is a different case: the plug-in can take it out of the pauses, but under the voice the two overlap, and above 60 % or so the voice begins to lose its air. There, 40 to 55 % and a little remaining background sound better than a clean pause with a hollow voice.
- **Sung vocals take less than speech.** A sung line's air, its reverb tails and the room it was recorded in are part of the sound; 25 to 55 % keeps them. Speech can take more because the gaps carry nothing you want.
- **Music takes the least.** On a stem or a finished mix, 15 to 35 % cleans the fades and the gaps without any risk to the music, which the plug-in never touches while it plays.

It glides. Turning the knob never clicks: the floor follows it smoothly, and moving to or from 0 % crossfades over 20 ms. REMOVAL can be automated, so you can clean a quiet intro harder than the chorus.

At 0 % the plug-in changes nothing: the output is the input, bit for bit, delayed by the latency. It still listens and learns, so the display keeps showing the background and you can bring REMOVAL up at any moment.

LISTEN

LISTEN plays **only what is being taken out**: the input minus the cleaned output. It is the quickest way to judge a setting. What you want to hear is the background alone, a steady hiss, a hum, the room, perhaps a faint ghost of the loudest syllables. What you don't want to hear is the voice: if you can make out words, consonants or the ring of a note, REMOVAL is too high for this recording. Turn it down until the words disappear from LISTEN, then switch it off and judge the voice itself.

Click LISTEN to switch it on; the button fills to show it is on. Click it again to go back to the cleaned sound. LISTEN is a listening aid, not a setting: it is never saved in a preset, it can't be automated, and it switches itself off when the host renders offline. With REMOVAL at 0 % it plays silence, because nothing is being taken out.

> **Tip:** With LISTEN on, the REMOVING read-out and the meter bridge still measure the cleaned output, not what you hear, so you can read how much is going while you listen to what it is.

Displays

The BACKGROUND display

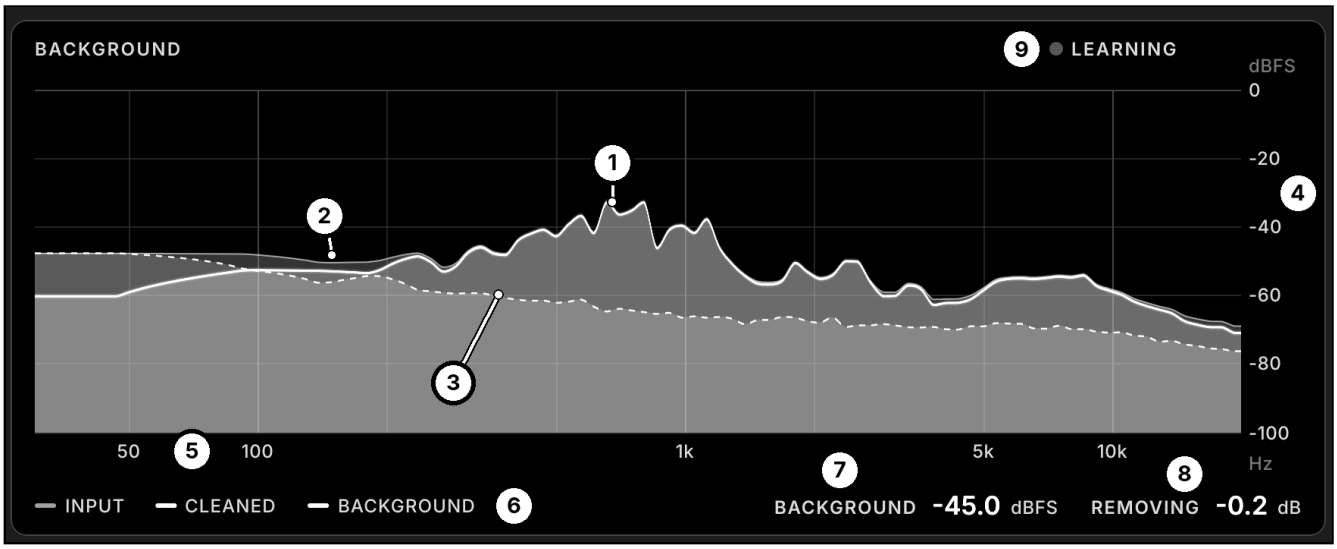


FIGURE 13 The BACKGROUND display while a voice plays: 1 the cleaned output (blue on screen, the filled curve), 2 the input (grey, a thin line along the top of the shaded band at the left), 3 the learnt background (orange on screen, dashed), 4 the level scale in dBFS, 5 the frequency scale in Hz, 6 the legend, the 7 BACKGROUND and 8 REMOVING read-outs, and 9 the LEARNING lamp, dark here because a word is playing.

The BACKGROUND display shows the sound from 30 Hz to 20 kHz, the lows on the left and the highs on the right, on a musical scale (each octave the same width), against a level scale 4 from 0 down to −100 dBFS. It draws three traces, named in the legend 6 :

- INPUT 2 is what goes into the plug-in: the voice and the background together. It is drawn in grey as a thin line with a faint fill. Wherever the plug-in takes something out, this line shows above the cleaned trace, and the band between the two is what went.
- CLEANED 1 is the output, drawn in blue on screen as the filled curve with the bright edge. Where the voice is, it sits on the input; where there is only background, it drops below it.
- BACKGROUND 3 is the picture of the background the plug-in has learnt, drawn in orange on screen as a **dashed** line. Without colour, it is the only dashed trace. The input runs along it in the pauses, and the voice rises above it while it plays.

The two read-outs at the bottom right update ten times a second:

- BACKGROUND 7 is the overall level of the learnt background, in dBFS: in the picture, −45.0 dBFS. It reads -- until the plug-in has heard its first pause.
- REMOVING 8 is how much quieter the output is than the input right now, in dB. While a voice plays it hovers near 0, because the voice is most of the sound; in the pauses it drops to the floor.

The LEARNING lamp (9) lights (green on screen) while the plug-in hears a pause and updates its picture. Watch it for a moment after you press play: it should light in the lead-in and in every gap between phrases. If it never lights, the plug-in has found no pause, and nothing will be cleaned (see *How it tells a pause from a word* and *Troubleshooting*).

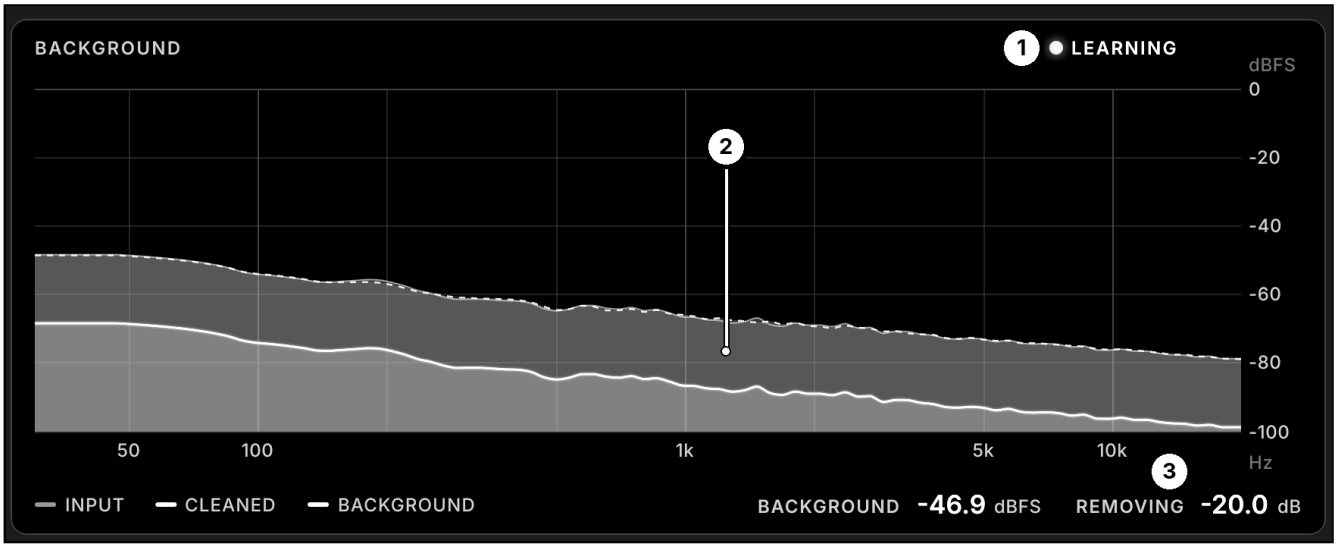


FIGURE 14 The same preset in a pause: (1) LEARNING is lit, the cleaned trace sits a steady (2) 20 dB under the input across the whole range (the floor at 50 %), and (3) REMOVING reads −20.0 dB. The dashed background runs exactly along the input, because the input *is* the background now.

Hover anywhere in the plot for a read-out of the frequency under the mouse and the three levels there, for example *1.2 kHz in −45.0 out −61.0 bg −60.0 dBFS*. A thin line marks the frequency you point at.

What to look for:

- **Before you set anything**, play a pause and watch the three traces meet: input, cleaned and background should line up, with the cleaned one a steady distance below. That distance is the floor.
- **Play the voice** and watch the cleaned trace follow the input over the hump of the voice and fall away from it only outside. If the cleaned trace sits visibly below the input *inside* the voice's range while words play, the plug-in is taking some of the voice: lower REMOVAL.
- **Hum** shows as narrow spikes at the far left of the background trace (50 or 60 Hz and its multiples). The plug-in learns them like any other background, so they disappear from the cleaned trace without any setting.
- **A background that changes** (a fan that speeds up, traffic that builds) shows as the dashed trace moving at each pause. Between pauses it holds still.

Meter bridge (IN | OUT)

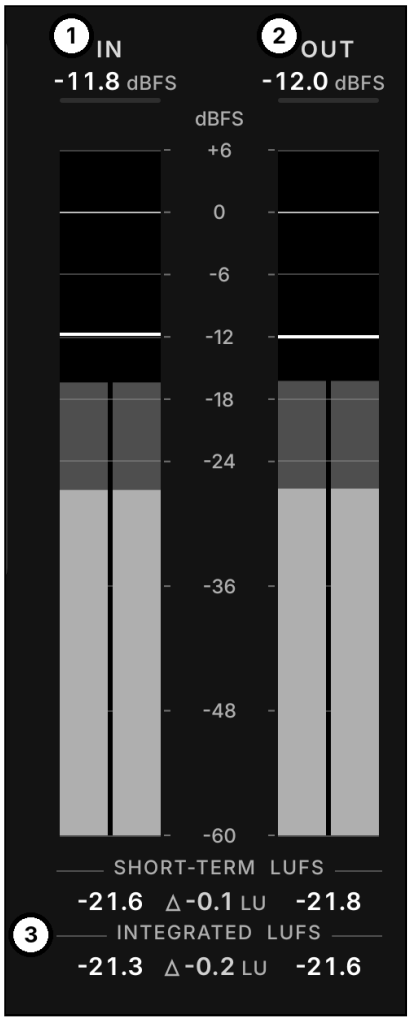


FIGURE 15 The meter bridge: 1 input and 2 output level, 3 short-term and integrated loudness with the OUT − IN difference.

VEIRA Clean Up's meter bridge is the standard one, measured on the plug-in's input and output. The Δ readings tell you how much quieter the recording got as a whole: a fraction of a decibel while a voice plays, since the voice is most of the sound (−0.1 LU in the picture), and the full floor in a long pause (−20 LU at 50 %). On a mono track, both bars of each meter show the same mono signal.

Reading the IN and OUT meters

The meter bridge measures what goes into VEIRA Clean Up and what comes out, in level and in loudness. The drawing shows the input and output columns; any extra column on your plugin's bridge is described above.

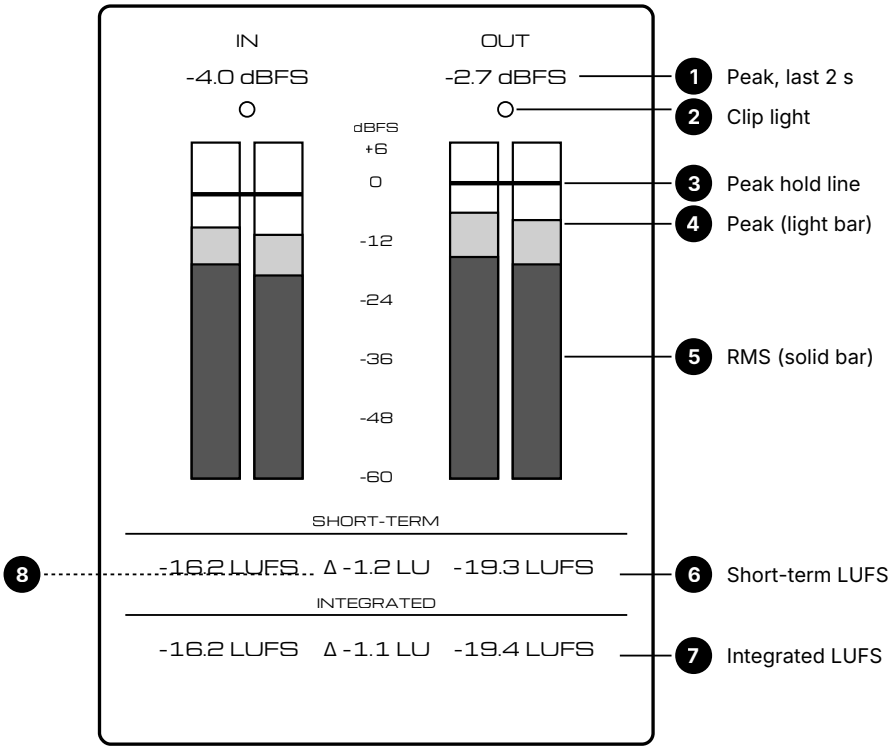


FIGURE 16 The meter bridge, left and right channel bars for input and output, with loudness below.

- 1 **Peak readout:** the highest sample peak of the last two seconds, in dBFS. It turns to a warning above 0 dBFS.
- 2 **Clip light:** lights on any sample at or over full scale. It stays lit until you click the meter.
- 3 **Hold line:** marks the recent peak on the bar.
- 4 **Peak bar** (the lighter part): the instantaneous sample peak of each channel.
- 5 **RMS bar** (the solid part): the average level over 300 ms, closer to how loud it sounds.
- 6 **SHORT-TERM:** loudness over the last 3 seconds, in LUFS, for the input (left) and the output (right).
- 7 **INTEGRATED:** loudness of everything since the measurement started, gated as the ITU-R BS.1770-4 and EBU R 128 standards require.
- 8 **Δ (delta):** OUT minus IN, in LU. It tells you how much louder (+) or quieter (−) VEIRA Clean Up made the signal, so you can match levels fairly when you compare with bypass. It reads [--] while either side is too quiet to measure (below −70 LUFS).

The scale: on screen the bars are coloured like a studio console: green below −12 dBFS, yellow from −12 to −3, orange from −3 to 0 and red above 0 dBFS.

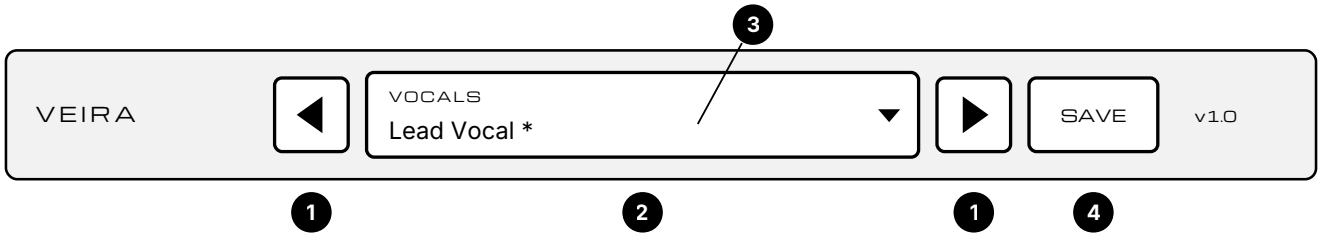
Clicks: click a level meter to reset its hold and clip light. Click the loudness readouts to restart the integrated measurement, for example at the start of a song.

Presets

VEIRA Clean Up comes with **27 factory presets** in 7 folders. Each one is a starting point made for a real job. Load one close to what you need, then adjust the first control its note mentions.

The preset browser

The preset browser sits in the header, at the top of the window.



The header: logo and plugin name, then the preset browser, then the version.

FIGURE 17 The preset browser.

- 1 ◀ ▶ step to the previous or next preset, through every folder in order.
- 2 **The display** shows the folder (small, above) and the preset's name. Click it to open the preset menu. Hover over it to read the preset's note: what it's for and what to adjust first.
- 3 **An asterisk (*)** after the name means you've changed something since the preset loaded.
- 4 SAVE stores your current settings as a user preset.

The preset menu

Clicking the display opens a menu with:

- one submenu per **factory folder**, in the order shown below;
- **User Presets**, with your own presets in their folders;
- **Save...**, **Delete** (only for user presets: factory presets can't be deleted), **Rescan** (re-reads the user folder, if you copied presets in by hand) and **Show User Preset Folder** (opens it in Finder or Explorer).

Saving your own presets

1. Set VEIRA Clean Up the way you like it.
2. Press SAVE. A dialog opens inside the plugin window.
3. Type a name, choose a folder or type a new folder name, and confirm.

Each user preset is one small .veirapreset file:

	Folder
macOS	~/Library/Audio/Presets/VEIRA/VEIRA Clean Up/<Folder>/
Windows	%APPDATA%\VEIRA\VEIRA Clean Up\Presets\<Folder>\

Copy these files to move your presets to another computer, or to back them up. Presets keep loading in future versions: any control added later simply starts at its default.

Presets in your DAW

The factory presets also appear in your DAW's own preset or program menu, named *Folder: Name* (for example *Vocals: Lead Vocal*). Whatever preset is selected is saved with your session and comes back when you reopen it, including the asterisk if you'd changed it.

The **Default ▸ Init** preset equals a fresh instance of the plugin: every control at its default.

The factory folders

Folder	Presets
Default	Init
Speech & Dialogue	Light Touch, Podcast Clean, Noisy Room, Phone Interview, Silence Between Lines
Vocals	Gentle Vocal, Studio Vocal, Bedroom Vocal, Exposed Vocal
Instruments	Acoustic Guitar, Piano, Amp Hum, Brass & Winds, Drum Overheads
Field & Location	Outdoor Interview, Room Tone, Nature Recording, Camera Mic
Music & Mixes	Stem Clean, Old Tape Transfer, Sample Clean, Whole Mix Safe
Restoration	Hiss, Hum & Buzz, Computer Fan, Maximum

Each preset is one REMOVAL setting for a situation, and its note (in the preset menu) says how much comes off and what to listen for. The folders are sorted by what is being cleaned:

- **Speech & Dialogue** runs from 20 % to 100 %: *Light Touch* (8 dB off) for a good recording that needs a polish, *Podcast Clean* (16 dB) for a home studio or an office, *Noisy Room* (24 dB), *Phone Interview* (30 dB) for a call recording, and *Silence Between Lines* (40 dB) for a radio spot or a dub whose lines must sit on silence.
- **Vocals** runs lower, 25 % to 70 %, because a sung line's air and reverb tails are part of the sound: *Gentle Vocal* (10 dB), *Studio Vocal* (16 dB, the everyday setting), *Bedroom Vocal* (22 dB, a fan and the street) and *Exposed Vocal* (28 dB) for a line that will sit alone in a breakdown.
- **Instruments** (20 % to 50 %) are gentle, so that an instrument's ring and decay survive: *Acoustic Guitar* (12 dB), *Piano* (10 dB), *Brass & Winds* (14 dB), *Drum Overheads* (8 dB) and *Amp Hum* (20 dB) for an amp that hums and buzzes between the notes.
- **Field & Location** (30 % to 80 %) is for recordings made outside the studio, where the background is louder: *Outdoor Interview* (22 dB), *Room Tone* (18 dB, enough to sit location dialogue in a mix without losing the room), *Nature Recording* (12 dB: only the recorder's hiss goes, the ambience is the sound) and *Camera Mic* (32 dB).
- **Music & Mixes** (15 % to 60 %) works in the gaps and fades of music and never touches it while it plays: *Stem Clean* (14 dB), *Old Tape Transfer* (18 dB), *Sample Clean* (24 dB, so a one-shot starts and ends on silence) and *Whole Mix Safe* (6 dB) for a finished mix.

- **Restoration** (50 % to 100 %) goes furthest: *Hiss* (20 dB), *Computer Fan* (24 dB), *Hum & Buzz* (26 dB) and *Maximum* (40 dB) for badly recorded material; back off if the voice turns watery.

Levels and backgrounds. The presets were measured on recordings whose background sits where the folder expects it: 33 dB under a spoken voice for Speech & Dialogue, 10 dB under it for Field & Location, 38 dB under a guitar for Instruments. On each, the sound kept its level within 0.3 dB (within 1 dB on the loud location backgrounds) and the background alone went down by exactly the floor. What a preset can't know is where *your* background sits, so treat the folders as starting points and check REMOVING and LISTEN on your own recording. None of the presets sets LISTEN.

Recipes

Step-by-step starting points. In each, let the plug-in hear a pause before you judge it, use LISTEN to check what is going, and compare with bypass now and then: the display shows what the plug-in decides, not how it sounds.

A podcast episode

1. Insert VEIRA Clean Up on each voice's track and load **Speech & Dialogue ▸ Podcast Clean**.
2. Play from the start of the recording, where there is usually a moment before the first word. LEARNING lights, and the BACKGROUND read-out shows the level of the room.
3. Listen to a passage with gaps. Raise REMOVAL until the fan and the room are gone from the gaps (40 to 60 % for most home recordings), then switch LISTEN on: you should hear the room, not the words.
4. If the gaps still carry a soft noise at 100 %, that is breath: the voice itself. Leave it, or put a gate after the plug-in for a hard silence.
5. Put your compressor and EQ *after* VEIRA Clean Up, so they work on the clean voice.

A voice-over that must sit on silence

1. Load **Speech & Dialogue ▸ Silence Between Lines** (100 %).
2. Play the lines and the gaps. With a quiet studio background the gaps go to nothing and the voice is untouched.
3. If the voice turns watery or loses its air, the background is too close under it for 40 dB: lower REMOVAL to 75 % (*Phone Interview*) or 60 %, and let the editor's fades do the rest.

A vocal recorded at home

1. Put VEIRA Clean Up first on the vocal track, before the compressor and the reverb or delay sends, and load **Vocals ▸ Bedroom Vocal** (55 %).
2. Play the verse. The computer fan and the street disappear from between the lines, and the sends stay clean because what reaches them is already clean.
3. Listen to the breaths and the tails of the long notes with LISTEN on: if you hear the singer in it, go down to **Studio Vocal** (40 %).
4. For a line that will stand alone in an intro or a breakdown, **Exposed Vocal** (70 %) so nothing hisses when the band stops.

An amp that hums

1. Load **Instruments ▸ Amp Hum** (50 %) on the guitar or bass track.
2. The plug-in needs to hear the amp alone: a moment before the first note or a gap between phrases. If the part never stops, record (or keep) a couple of seconds of the amp idling at the start of the clip.
3. In the gaps the hum and the buzz go down 20 dB; while the notes ring, only the hum between their harmonics goes. The notes' decay is kept, because the plug-in waits for a real pause before it calls anything background.

An interview recorded outside

1. Load **Field & Location ▸ Outdoor Interview** (55 %).
2. A street or a windy field is a loud background, often louder than −40 dBFS. The plug-in learns it at the first gap *after* the first words, so start playback a few seconds before the clip you're judging, or let the first answer go by once.
3. Expect a cleaner, closer voice rather than silence: with the background only 10 dB under the voice, taking all of it would hollow the voice. If the voice thins, go to **Room Tone** (45 %) and keep a little of the place in the sound.

A sample or an old tape transfer

1. Load **Music & Mixes ▸ Sample Clean** (60 %) on a one-shot, or **Old Tape Transfer** (45 %) on a transfer.
2. Make sure the clip starts with its noise: the hiss before the first note is what teaches the plug-in. Trim *after* cleaning, not before.
3. The noise before and after the sound goes down 24 dB (18 for the transfer) and the music itself is untouched while it plays. For a finished mix, use **Whole Mix Safe** (15 %): it only works in the fades and the gaps.

Specifications

Type	Background sound remover: takes room tone, hiss, hum, fans and traffic out from under a voice or an instrument; the background is learnt automatically in the pauses
Channels	Mono and stereo. Stereo is analysed and cleaned as one (the same gain for both sides, from their summed power), so the stereo image never moves
Latency	1,024 samples at 44.1 and 48 kHz (23 and 21 ms), 2,048 at 88.2 and 96 kHz, 4,096 at 176.4 and 192 kHz: one analysis frame. Constant, reported to the host for delay compensation
Transparency	REMOVAL 0 %, and bypass: the output is the input delayed by the latency, bit for bit. Moves to and from 0 %, and LISTEN, crossfade over 20 ms
Analysis	Overlapping frames of 1,024 samples (21 ms at 48 kHz), one every 256 samples (5.3 ms), about 500 frequencies per frame
Learning	A frame counts as a pause when it is within 6 dB of the quietest frame of the last 3 to 6 s and either 12 dB under the loudest or under −40 dBFS (100 Hz to 8 kHz). In a pause the picture follows the background within 0.1 s, frequency by frequency; between pauses it is kept. Sound that never pauses is left unchanged
REMOVAL	0 to 100 % in 1 % steps (default 50 %). The floor: 0.4 dB per percent, −10 dB at 25 %, −20 dB at 50 %, −40 dB at 100 %. In a pause the background goes down by exactly the floor; under the voice only what is clearly background
LISTEN	On, off (default off): plays the input minus the cleaned output. Not saved in presets, not automatable, off when the host renders offline
Measured	A spoken voice at −12 dBFS peak over hiss and hum at −45 dBFS, 48 kHz, REMOVAL 50 %: the gaps between words down 16.5 dB, the voice's level within 0.4 dB; the background alone down 10.0, 20.0 and 40.0 dB at 25, 50 and 100 %
BACKGROUND display	Input, cleaned output and learnt background, 30 Hz to 20 kHz on a musical scale, 0 to −100 dBFS; BACKGROUND (dBFS) and REMOVING (dB) read-outs; LEARNING lamp; hover read-out per frequency
Metering	IN and OUT meters (sample peak, RMS); short-term and integrated loudness per ITU-R BS.1770-4 / EBU R 128
Presets	27 factory presets in 7 folders, plus your own
Window	1160 × 720 points; resizable from 75 % to 200 %, saved with the session
Formats	VST3, Audio Units, standalone; AAX once signed
Plugin ID	Manufacturer Veir, plugin Vcln

Troubleshooting and questions

The sound drops out for a moment every 30 seconds. The plugin isn't activated on this computer. Open VEIRA Manager and sign in; the strip at the bottom of the plugin's window says exactly what's wrong. See *Getting started* ▸ *Activation and the licence*.

I don't see VEIRA Clean Up in my DAW. Make sure VEIRA Manager shows SHOW (installed) for it, then rescan plugins in your DAW. In Logic Pro, open the Plug-in Manager and choose *Reset & Rescan Selection* if needed. Look under the manufacturer **VEIRA**.

Nothing happens: REMOVING reads 0.0 and BACKGROUND reads --. The plug-in hasn't heard a pause yet, so it has no background to remove. Play from a point where the recording has a second of background alone and watch LEARNING light. If it never lights, the sound has no pauses: a held note, a strummed guitar or a dense mix never drops 12 dB under its loudest moment, and the plug-in leaves it alone by design.

The first second of my clip isn't cleaned. With a loud background (over −40 dBFS) the plug-in only recognises a pause once it has heard something louder, so the first phrase passes as recorded and the cleaning begins at the first gap after it. Start playback or the bounce a few seconds earlier, or leave a gap after the first sentence.

The voice sounds watery, thin or robotic. REMOVAL is too high for this recording. Switch LISTEN on: if you hear words, lower REMOVAL until they disappear, usually to 30 to 60 %. A background close under the voice (10 dB under it) can't be taken out of the voice without harming it; accept a little of it under the words.

The background is still there under the voice. Under the voice the plug-in takes only what is clearly background, to keep the voice whole; raise REMOVAL a little. If the background is nearly as loud as the voice, no setting removes it from under the words cleanly: a closer microphone next time is worth more than any plug-in.

The gaps between words aren't silent even at 100 %. What's left is breath, and the tails of words: they are the voice, and the plug-in keeps them. For hard silence between lines, follow VEIRA Clean Up with a gate, or edit the gaps.

The whole recording got quieter. Very quiet material, under −40 dBFS and hardly changing in level, is taken for background. Raise the level going into the plug-in (clip gain, or a gain plug-in before it) so the sound sits clearly above −40 dBFS.

The hum hasn't gone. Hum is learnt like any other background, but only in the pauses: make sure LEARNING has lit at least once since the hum began. If the hum changes (a dimmer, a motor), the picture catches up at the next pause. *Restoration* ▸ *Hum & Buzz* takes 26 dB off it and its harmonics.

The window is too small (or too big) on my screen. Drag its bottom-right corner. It scales from 75 % to 200 % and keeps its proportions; the size is saved with your session.

How do I type an exact value? Right-click REMOVAL (Control-click on a Mac), type the number and press Return.

Do my sessions keep my settings? Yes. REMOVAL, LISTEN, the selected preset and the window size are saved with your session. REMOVAL can be automated; LISTEN can't, presets never set it, and an offline bounce switches it off by itself (in a real-time bounce it stays as you left it).

Contact: support@veiraaudio.com. Mention your computer (Mac or Windows and its version), your DAW and version, and what you see.



VEIRA Clean Up · User Manual · Version 0.2.1

Support: **support@veiraaudio.com** · Web: **veiraaudio.com**

© 2026 VEIRA. Proudly made in the Dominican Republic. VST is a trademark of Steinberg Media Technologies GmbH. Audio Units, Logic Pro and macOS are trademarks of Apple Inc. Pro Tools and AAX are trademarks of Avid Technology, Inc. All other trademarks belong to their owners and are used only to describe compatibility.