



FORESIGHT

OFFLINE COMPRESSION PREVIEWER

User Manual

Renegade Sound Labs

Version 0.1.0 · Free Plug-in · macOS and Windows

What Foresight is

Foresight shows you what a compressor *would* do to your audio, before you commit to compressing it.

It captures up to sixty seconds of whatever is passing through it, then runs a full compressor simulation over that capture and draws the result: the input, the compressed output, and the gain reduction envelope. Move a knob and the whole capture is re-simulated and redrawn. You are looking at the actual consequence of a setting on your actual material, not a meter reacting in real time.

Foresight never processes your audio. The signal passes through completely untouched. There is no latency, no oversampling, and no delay compensation to worry about. You can leave it anywhere in a chain and it changes nothing. Every compressed waveform it draws is a simulation.

Why that is useful

A gain reduction meter tells you what just happened, for an instant, while the music moves past. Foresight holds a passage still so you can study it: where the compressor grabs, how long it takes to let go, which transients survive and which get flattened. You can compare two settings on the same performance instead of on two different bars.

Getting started

1. Insert Foresight on a track or bus.
2. Press **CAPTURE**. The button turns red and reads **CANCEL** while it waits.
3. Roll the transport. Recording starts as soon as the DAW plays and the readout counts up.
4. Press **STOP** when you have enough, or let it run to the full sixty seconds.
5. Move the knobs. The display redraws continuously as you drag.

Installation

Foresight runs on macOS 11 or newer and on Windows 10 64-bit or newer. Quit every DAW before installing.

MACOS Installer

Open `RSL-Foresight-v0.1.0.pkg` and follow the prompts. The installer is signed and notarized, so it opens with no Gatekeeper warning. You can choose the Audio Unit, VST3, Standalone and manual separately.

Audio Unit	/Library/Audio/Plug-Ins/Components/
VST3	/Library/Audio/Plug-Ins/VST3/
Standalone	/Applications/, with the uninstaller alongside it
Manual	/Library/Application Support/Renegade Sound Labs/

An uninstaller is included and lands in your Applications folder. It removes every component and asks whether to keep your saved presets.

WINDOWS Installer

Run `RSL-Foresight-v0.1.0-Windows.exe` and follow the prompts.

SmartScreen will warn you on first run. The Windows installer is not code signed yet, so Windows shows a blue “Windows protected your PC” panel. Choose **More info**, then **Run anyway**. This is expected, and it will stop once the installer is signed.

VST3	C:\Program Files\Common Files\VST3\
Standalone	C:\Program Files\Renegade Sound Labs\

There is no Audio Unit on Windows. AU is a macOS format, so the Windows installer offers VST3 and Standalone only.

After installing

Reopen your DAW and rescan plug-ins if it does not find Foresight automatically. Logic Pro scans Audio Units at launch, so quitting and reopening it is enough.

The display



The full interface after a capture, with a suggestion applied and a point pinned on the display.

Capture states

READOUT	MEANING
READY	Nothing captured. Press CAPTURE to arm.
ARMED	Waiting for the transport to roll. Nothing is being recorded yet.
RECORDING	Writing to the capture buffer, counting up to 60 s.
CAPTURED	Buffer holds audio. Length and sample rate are shown.

Pressing the button again at any point cancels or stops. RECAPTURE discards the current capture and arms a fresh one.

Reading the display



Gain reduction across the top, input and simulated output below, with the threshold drawn as dashed lines and a readout pinned to a point in the audio.

Gain reduction strip

The upper pane is the gain reduction envelope, hanging down from zero. Gridlines mark every 6 dB and the strip rescales to whatever the current settings actually reach. The deepest reduction anywhere in the capture is printed at the top right as PEAK.

Waveform pane

TRACE	WHAT IT IS
Silver	The captured input, exactly as it arrived.
Amber	The simulated compressor output.

Where the two coincide, nothing is happening: the amber simply covers the silver. Where the compressor is working, the silver shoulders show above and below the amber. That gap is the gain reduction. The dashed horizontal lines mark the threshold, mirrored above and below the centre.

Reading a point

Hover anywhere on the display for a crosshair and a readout box giving the time position and the input level, output level, and gain reduction at that instant. A ring marks the point on the output waveform and a dot marks it on the gain reduction envelope.


Click to pin a point. Pinned points stay put while you move the knobs, each carrying its gain reduction figure, so you can watch a specific transient respond as you dial. Click a pin again to remove it, or use CLEAR PINS.

Zooming

Scroll	Zoom in and out around the pointer.
Drag	Pan through the capture.
Double-click	Fit the whole capture.
FIT	Same as double-click.

Zooming genuinely re-simulates at higher resolution rather than magnifying pixels, so detail sharpens as you go in. Pins survive zooming. The visible time range and zoom factor are shown at the bottom of the pane.

Controls



The control row. Double-click any knob to return it to its default; hold Shift while dragging for fine adjustment.

CONTROL	RANGE	WHAT IT DOES
Threshold	-60 to 0 dB	The level above which compression begins. Drawn on the display as the dashed lines.
Ratio	1:1 to 20:1	How much of the level above the threshold is removed. 1:1 is no compression.
Knee	0 to 24 dB	How gradually compression engages. 0 is a hard knee: nothing, then full ratio. A wider knee starts easing in below the threshold and reaches full ratio above it.
Attack	0.1 to 200 ms	How quickly gain reduction builds when the signal crosses the threshold. Short attacks catch transients; longer ones let them through.
Release	5 to 2000 ms	How quickly gain reduction lets go afterwards. Watch the recovery curve on the strip: too short pumps, too long never recovers between hits.
Lookahead	0 to 20 ms	Lets the detector see ahead of the audio, so reduction is already underway when the transient lands. Costs nothing here, because Foresight is offline and has no latency to compensate.
SC HPF	OFF, 20 to 500 Hz	High-passes the detector only, so bass energy stops driving the compression. The audio itself is unaffected. Reads OFF at the bottom of its range.
Makeup	0 to 24 dB	Gain applied after compression, to bring the level back up.

Detector: AUTO / PEAK / RMS

PEAK follows the instantaneous level and reacts to every transient. RMS averages over a short window and ignores brief peaks, the way a bus compressor does.

AUTO lets SUGGEST choose the detector to suit the source you picked. Clicking PEAK or RMS pins your choice: suggestions from then on build around it rather than replacing it, and the threshold is re-solved so the goal is still met with the detector you asked for. Press AUTO to hand the choice back. PEAK and RMS always

show which detector is actually running, so AUTO lit alongside RMS means “chosen for you, and it picked RMS”.

A/B

Two independent setting slots. Switch between them to compare two treatments of the same capture. COPY copies the current settings into the other slot, which is the usual way to start: copy, then vary one thing. Both slots are saved with your session.

Suggest

Tell Foresight what it is listening to and what you want from it, and it fills in all nine controls against your actual capture.

SOURCE	WHAT IT SETS	
Vocal	Medium attack, soft knee, fast enough for consonants without dulling them.	
Drums	Fast attack, hard knee, sidechain filter so kick energy does not duck the kit.	
Bass	Slow attack, so it follows the note rather than chasing each cycle.	
Guitar / Keys	Middle ground, soft knee.	
Room / Overheads	RMS detection and a long release, so decay does not pump.	
Full Mix	Bus settings: RMS, high-passed sidechain, slow timings, lookahead.	
GOAL	RATIO	TARGET PEAK GAIN REDUCTION
Glue	2.0:1	4 dB. Gentle, always working, holds parts together.
Balance	3.0:1	7 dB. Evens the performance so nothing disappears.
Control	4.5:1	10 dB. Firm, catches peaks and leaves the rest alone.

The threshold is not calculated from a rule of thumb. Foresight searches for it: it runs the compressor over your capture repeatedly, adjusting the threshold until the measured peak gain reduction matches the goal. That is why the same goal behaves consistently across quiet and loud material, and why the number it targets is the same PEAK figure shown on the gain reduction strip. You can check its work.

Makeup is measured the same way: a final pass finds the compressed peak and sets makeup to restore the original level, so switching a suggestion in and out is a fair comparison rather than a loudness contest.

After a suggestion, the line to the right of the button reports what actually happened: which detector was chosen or kept, the measured peak gain reduction, and the makeup applied.

Suggestions are starting points, not answers. They put you in the right area on this material so you can hear and see what to change. Nothing stops you moving every control afterwards.

Working with it

Capture something representative

Capture the loudest, busiest section you care about, not an intro. The compressor responds to what is there, and a threshold placed against a quiet passage will be wrong everywhere else.

Watch the recovery, not just the depth

Most compression problems are timing, not amount. Pin a point just after a transient and vary release: you are looking for the envelope to return to zero before the next hit, without snapping back so fast it pumps.

Use the shoulders

The silver above the amber is the reduction, drawn on the waveform itself. If a transient you wanted to keep has been flattened, you will see its silver peak standing alone above a blunt amber edge. Lengthen the attack.

Then go and set your compressor

Foresight does not compress. Once you know the settings you want, dial them into the compressor you actually intend to use. The model is a clean, transparent feed-forward design; a compressor with its own character will not sound identical, but the threshold, ratio, and timing relationships carry across.

Specifications

Formats	macOS: Audio Unit, VST3, Standalone · Windows: VST3, Standalone
System	macOS 11 or later, or Windows 10 64-bit or later
Architecture	macOS: Universal 2, native on Apple Silicon and Intel · Windows: 64-bit Intel or AMD
Channels	Mono or stereo
Latency	Zero. No oversampling, no delay reporting.
Audio path	Bit-transparent pass-through. The plug-in never alters the signal.
Capture	60 seconds, stereo, at the host sample rate
Model	Feed-forward, log domain, per sample. Peak or RMS detection, soft or hard knee, stereo linked on the louder channel, optional detector high-pass and lookahead.
Licence	Free. No key, no time limit, no watermark.

Support

If something looks or behaves wrong, please tell us. It helps to include:

- Operating system and version, and whether you are on macOS or Windows
- DAW and plug-in format (AU, VST3, or Standalone)
- Sample rate and buffer size
- The control settings, and what you expected to see
- Steps to reproduce

hello@renegadesoundlabs.com · renegadesoundlabs.com