

Clarity Vx Pro

Real-Time Noise Reduction for Voice
Powered by Waves Neural Networks®



User Guide

Contents

INTRODUCTION	3
QUICK START	4
Understanding the Mix	4
Interface	5
Step-by-Step	6
Basic Controls	6
Advanced Controls	6
Advanced Controls—Graphic Section	6
Advanced Controls—Bottom Panel	7
CONTROLS	8
Bottom Panel	8
Analysis Mode	8
Sensitivity	9
Reflections	9
Neural Network Selector	10
Reset	11
AUTO RESET	11
Main Control Knob: Voice vs. Ambience	11
Process Amount	12
Ambience Gate Section	12
Output Section	13
Graphic Display	14
Controlling the process amounts per band	15
Band Process Controls	16
Keyboard Shortcuts and Precision Controls	17
WaveSystem Toolbar	18
CPU	18
Error Message	18



Introduction

Thank you for choosing Waves! To get the most out of your new Waves plugin, please take a moment to read this user guide. Installing software and managing your licenses require a free Waves account. Sign up at www.waves.com. With a Waves account you can keep track of your products, renew your Waves Update Plan, participate in bonus programs, and keep up to date with other important information.

We suggest that you become familiar with the Waves Support pages, www.waves.com/support, where you will find technical articles about installation, troubleshooting, specifications, contact information, and more.

Neural Networks

A neural network is, in essence, a learning machine. It uses algorithms inspired by the brain to store information and make choices. Before a neural network can do anything useful, such as distinguishing ambience from voice and acting on that information, it must analyze millions of relevant samples. As it learns more and more from these diverse samples, the neural network creates a set of associations. When an input signal is sent through the network, it is analyzed, and a set of weights and biases are created. This is the information that is used to control processing. The more samples it learns from, the better its decision-making skills. The choice of material used to teach neural network greatly influences its decisions.

Waves Clarity Vx Pro

Clarity Vx Pro uses a deep neural network to adaptively isolate even the most complicated background sounds and separate them from the voice. It can give you voice without ambience, ambience without voice (or anything in between), and provide powerful background noise reduction with almost no impact on the voice. For basic use, a single knob controls the mix of the fully separated voice and the ambience.

This is often all you need. Advanced controls can be used to adjust the process amount and gain in six adjustable bands. You can select how source material is analyzed, adjust how unvoiced regions are treated, hear the delta (DIFF), and control the stereo image of the output.

Clarity Vx Pro has mono and stereo components. It can be used in real time as a channel insert, or as an offline processor.



Quick Start

This quick start will teach you how to isolate ambience from voice sound with minimum impact on either. Follow these steps and you'll achieve a high-quality separation of voice from ambience. Refer to the rest of this user guide to learn more about Clarity Vx Pro controls and how to use them for best results.

Understanding the Mix

The key to superior results lies in controlling the mix between the input signal and the output of the neural network. There are three controls that govern this mix; it is important to understand how each control works and how they interact.

MAIN CONTROL KNOB (AMBIENCE/VOICE)

The Main Control knob sets the range of how much ambience will be separated from the voice. Let's say you are using Clarity Vx Pro to remove ambience and you set the Main Control Knob to 80%. No matter how the other controls are set, the voice processing mix will never exceed 80%. For maximum flexibility, we suggest that you initially set this control to 100%.

BAND PROCESS FADERS

Each band in the Frequency Analyzer display has a draggable fader bar. This lets you control how the signal is processed by the neural network, per band. It lets you *reduce* processing per band, not *add* it. At the topmost setting, you are hearing 100% of the mix of the Main Control Knob in that band. As you lower the fader, you hear increasingly more of that band's ambience (i.e., you are moving back toward the bypass/input signal).

PROCESS AMOUNT

This is an intelligent VCA that links all the Band Process faders so that you can move them together. This lets you set up your band-by-band processing and then move them as a group.

BOTTOM LINE

The individual Band Process faders, which are linked together by the Process Amount control, determine the amount of processing applied to each band. The Main Control knob is a mix between the input signal and the neural network output, as defined by the other controls.



Interface

(Basic and Advanced controls shown.)



BASIC CONTROLS

- 1 Main Control: voice vs. ambience
- 2 Neural Network Section:
NN Selector, Reset, Auto-Reset, ECO mode indicator
- 3 Output Section: Level and Width
- 4 Advanced Controls Show/Hide

ADVANCED CONTROLS

- 5 Input Signal settings: Analysis and Sensitivity
- 6 Reflections: restores natural reflections
- 7 Process Amount Master: controls all bands together
- 8 Ambience Gate: amount of ambience left in unvoiced parts
Mix, Threshold, Attack, and Release
- 9 Graphic Frequency Analyzer View Options
- 10 Band Controls:
Solo, Bypass, Delta, and Gain
- 11 Adjustable crossover
- 12 Band Process Amount
- 13 Signal Display:
Input, Gain Reduction
- 14 Limiter

Tooltips



Hover over a control for a brief description of its function. The setting of the control is shown in a nearby value box.

Turn tooltips on or off in the WaveSystem Toolbar drop-down menu, next to the Save button.



Step-by-Step

Basic Controls

- Start with the **Main Control** knob. It sets the range of processing. Turn the knob to the right for more voice, turn it to the left for more ambience. For basic processing, find a setting that isolates the desired amount of voice or ambience. Some Advanced controls are limited by this setting, so if you plan to use the Advanced controls for band-by-band adjustments, we suggest that you begin with the maximum Main Control knob setting of 100.

- There are currently three neural networks. Each has a different effect on the balance of the voice signal. Use the **Neural Network** menu to select one.

The neural networks used by Clarity Vx Pro are each trained differently. Its training determines what a neural network is best suited for. This could be isolating voice from an especially harsh ambience or focusing on a foreground voice while eliminating other voices (or keeping them), or many other instances. Use the tooltip descriptions on the Neural Network menu to find the neural network that best suits your needs. Try more than one neural network since the results are content-dependent.

- Consult the Meter to set the output levels using the output **Faders** and adjust the size of the output stereo image with the **Width** control.

Advanced Controls

Use the Advanced Controls to adjust the processing amount and mix across six bands, enhance the sidechain for improved analysis, and control the post-processing reflections. The Advanced Controls view/hide button is beneath the Main Control knob.

Advanced Controls—Graphic Section

- Focus first on the six Band Process faders (these are the horizontal lines that are initially at the top of the graphic display). When a Band Process fader is at its highest setting, you are hearing the sound of the Main Control knob, at that frequency band. Drag a fader downward to offset the processing. The lower you drag the line, the less processing you'll hear in that band. Note that these band settings are dependent on the Main Control settings; a fader can reduce processing in a band, but not increase it. If you are not achieving enough ambience reduction with the Band Process faders, increase the value of the Main Control knob.



- Drag a **Crossover** left or right to adjust the range of a band. You can **Solo** or **Bypass** one or more bands, and you can monitor the difference between the input and processed signals with the **Delta** control. Use keyboard shortcuts and precision controls to enhance your productivity (see p. 16).
- Once you're satisfied with the wet/dry process amount of a band, use the **Pure Gain** control to raise or lower the level of the band with respect to the other bands.

Advanced Controls—Bottom Panel

- Drag **Process Amount Master** to group the faders of all six bands and move them together. This lets you adjust each band and then move them as a group while maintaining the ratios between them.
- **Analysis** and **Sensitivity**. Analysis controls how a stereo signal is analyzed before neural network processing: either two separate stereo tracks or a sum of the two. Sensitivity dynamically changes the input signal-to-noise ratio for improved analysis under certain circumstances. These settings do not change the audio signal, they only change the sidechain going into the neural network.
- Use the **Reflections** control to change the release time of the processor. This restores the natural reflections of the input signal, which may have been reduced during processing. Higher settings will restore more reflections, but may also bring back some noise during the reflections tail if the input signal has loud ambience.
- The **Ambience Gate** allows you to keep some ambience in the voiced section but attenuate or completely remove it during unvoiced sections. You cannot use this control to *increase* the level of ambience in unvoiced sections.
 - **Mix** controls the amount of ambience; the higher the setting, the more ambience is *removed* between voices.
 - The **Threshold** sliders set the level at which ambience reduction will begin and end.
 - **Attack** sets the residual ambience reduction speed.
 - **Release** sets the residual ambience recovery speed.
- If you notice the clipping LED, turn on the **Peak Limiter (LMT)** and safeguard your output from unwanted clipping artefacts. The Gain Reduction meter and Peak Gain Reduction value will be displayed.

All Clarity Vx Pro processing is carried out on your computer. It is not sent to a cloud for processing, so all operations can be carried out without internet connection.



Controls

This section describes all Clarity Vx Pro controls, so you'll need to be in Advanced Controls view to follow along. The controls are discussed in two parts: the Bottom Panel and the Graphic Display.

Bottom Panel

Before a signal can be separated into voice and ambience, it first must be analyzed. The Analysis and Sensitivity controls are used to optimize the input signal so that it is best suited for neural network processing. These controls do not affect the actual signal, just the detection. Compare this to a sidechain, which can equalize, compress, and gate the signal to provide better information to the dynamics processor—without changing the input signal.

Analysis Mode (stereo component only)

In the stereo component, left and right channels can be analyzed independently to create a unique set of weights and biases for each channel, or they can be summed together for a single analysis. Weights and biases provide instructions for neural network processing; two analyses mean that the neural network runs two processes.

- **Single** sums the left and right input channels before analysis. If the ambience and voice are similar on both channels, you don't need to analyze or process the input signal twice. In Single mode, both channels are processed using the same analysis data. Single mode consumes considerably less CPU than Double.
- **Double** mode is used when the left and right channels are significantly different (i.e., very different ambience or voice). Double mode analyses the left and right channels independently and processes them separately. This results in more precise processing for each channel, but it requires substantially more CPU. You may need to use the Width control to further adjust the stereo image.

Note that in some circumstances, Single will not yield the same degree of voice/ambience separation as Double. When in doubt, experiment with both since this is very content-dependent.



Sensitivity

Sensitivity changes the signal-to-noise ratio of the input signal to aid analysis. Normal sensitivity is usually the appropriate choice. But in cases where the signal-to-noise ratio is very low, try the High setting. This enhances the sidechain dynamics, which can improve the separation of certain noises and ambiences, especially high-end hiss and similar High-Frequency profile ambiences.



Remember, the Analysis and Sensitivity settings do not affect the input and output signals. They affect *only* the analysis. In a stereo signal, both left and right signals are always processed, and the sound image is not altered.

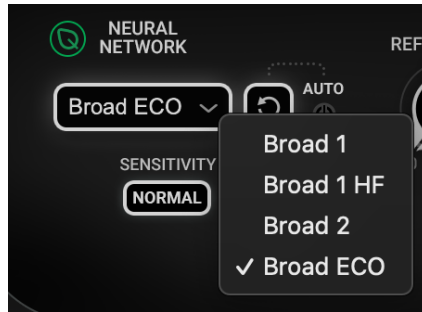
Reflections

Reflections controls the release times of the neural network process, after analysis. Low settings can yield a sound that is very accurate, but the voice's space and natural reflections may be reduced or cut into. A high Reflections setting will maintain the natural tail of the original signal, but it may bring back some of the original ambience. Settings above 150 ms can result in a sound that is pleasing, but there's risk of smearing, like a plate reverb. Appropriate settings typically lie between 30 ms and 200 ms.

Range: 0 ms to 1000 ms



Neural Network Selector



There are four Clarity Vx Pro neural networks. They differ in the statistics, signal-to-noise ratios (which were introduced to it during training), and weights and biases (which they apply during processing). Choose a network based on the input source, as well as the voice focus that you are seeking.

Broad 1	Isolate primary and secondary voices from ambience. This preserves the main voice, as well as some secondary voices, and isolates them from the ambience.
Broad 1 HF	Separate primary and secondary voices from ambience. This is similar to Broad 1, but with a greater high-frequency focus. This is not a shelf; it reflects different training biases.
Broad 2	Isolate the primary voice from secondary voices and ambience. This choice focuses on the main speaker/singer, separating it from the ambience and the secondary sounds (e.g., background voices, noises such as vacuum cleaners, traffic, or construction sites).
Broad ECO	This neutral network is optimized for reducing noise in dialog (spoken word/speech). It is a CPU-efficient neural network that lets you use more than twice the number of instances. It also provides for shorter offline bounces and reduced render times. When the ECO mode is selected, a small green leaf lights up. Broad ECO is not intended for reduction or removal of noise from sung vocals.

When using the Broad ECO neural network:

If you increase clip gains to the point at which Clarity's input goes beyond 0dB, we strongly recommend that you switch to the High Sensitivity mode.

Moreover, clip gains that cause Clarity's input to exceed the digital clip point (0dB) by 10 dB or greater will not be processed properly by the ECO neural network; in such cases, you should use the Broad 1 or Broad 2 neural network.



Reset

Clarity Vx Pro features deep neural networks that are capable of adaptive and continuous processing. As a neural network processes audio, it accumulates a “history” that influences its future decisions. However, to ensure reproducibility from one playback to another, Clarity resets the neural network every time you stop.


When the DAW plays a region in a loop, the neural network is constantly adapting, thus improving its performance with each pass. In certain cases, this can result in progressively better results with each loop cycle. So, if you encounter a particularly problematic section of your program material, we suggest that you loop-play the segment and loop-record several passes, until you reach satisfactory results. In most of these cases, you won’t need more than three passes to achieve optimal results.

On very rare occasions, playing a loop or moving the play head/cursor to a new starting point will cause degradation. If this occurs, click Reset to clear the neural network history and start anew. It is best to reset the neural network during an unvoiced region or silence. This can be automated.

AUTO RESET (DEFAULT: OFF)

Turning Auto-Reset on will automatically Reset the neural network every time it detects **200 ms of inactivity** (levels below -144dB). This can be beneficial in real-time use, when one instance of the plugin is expected to treat **several different voices or noise profiles** that are separated by at least 200 ms of silence.

Main Control Knob: Voice vs. Ambience



This sets the mix between voice and ambience, and therefore the overall amount of processing.


As you turn the knob clockwise, you will gradually hear less ambience, and the control will become more purple. At the highest setting, only voice is heard. Turn in the other direction to gradually hear less voice; the knob becomes increasingly blue and more intense. These colors are reflected in the graphic display to illustrate the voice/ambience balance.

The Main Control knob is after the Band Process faders in the Clarity Vx Pro signal flow. Hover over or touch the control knob to display its current value.

Range: -100 to +100




Process Amount



This links all the individual Band Process faders in the Graphic display. When the Process Amount control is moved up or down, the process amounts for all the bands will move together, maintaining the ratios between them. The controls converge as they approach 0 or 100 and will return to their previous positions as the Process Amount control moves them back to the center. You never lose your starting point.

Click the button below to disable all Band Process faders and revert to the Main Control's settings. This control cannot be automated, but you can automate each of the six Band Process faders. Band-based mixing is discussed in the next section. Range: 0 to 100

Ambience Gate Section



This section controls how much ambience is left in the gaps between the voiced parts. This affects the signal after the Main Control knob, so ambience between voiced sections can be reduced, but not increased.

The Ambience Gate is active when ambience is being reduced by less than 100%. This occurs when the Main Knob is set between 0.1% and 99.9% or when any of the Process Amount faders are set to less than 100%.

The Ambience Gate is not active when removing voice from ambience (Main Knob settings from -0.1% to -100%).

The Ambience Gate section consists of these controls:

Ambience Gate Mix sets the amount of ambience attenuation within non-voiced areas. So, as you increase the setting, more ambience is removed. Range: 0 to 100

The **Threshold Open** and **Threshold Close** markers set the levels between which ambience attenuation will occur during unvoiced sections. Range: -90 dB to 0 dB; Threshold Open reset value: -32 dB; Threshold Close reset value: -37 dB

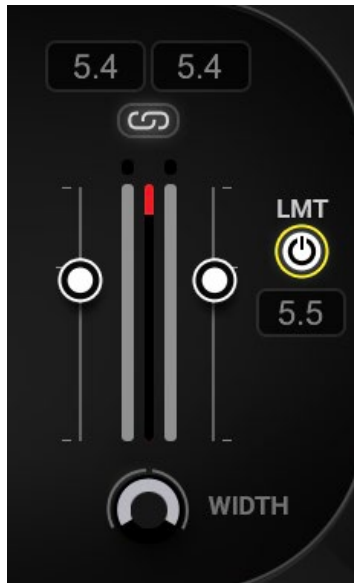
Attack sets the residual ambience reduction speed (in milliseconds). This is the speed at which the ambience will be reduced during unvoiced sections.



Release sets the residual ambience recovery speed (in milliseconds). This is the speed at which the ambience will return at the end of unvoiced sections.

Attack and Release Ranges: 0 ms to 500 ms; Reset values: 15 ms

Output Section



Full-scale meters (stereo or mono) indicate the plugin output. Peak hold values are shown above the meters. A red clip LED indicates a level of 0 dBFS. Click on the meter to clear peak hold and clip indicators.

Faders can move separately or together in the stereo component. Click the Link button above the meter to link them. If there is an offset between the two faders, it will be maintained when linked faders are moved.

Range: -inf to +24 dB

Use the **Width** control to adjust the size of the stereo image.
Range: 100 (original stereo width) to 0 (mono)

Engaging the **LMT** button will turn on the Output Limiter. This brick-wall limiter is applied at the absolute end of the signal chain and can be used to protect from harsh clipping artefacts and to ensure the set output level is not exceeded. When the limiter is active, a corresponding **Gain Reduction Meter** and **Gain Reduction Peak Hold value** are displayed, which allows for smarter, more informed decisions.

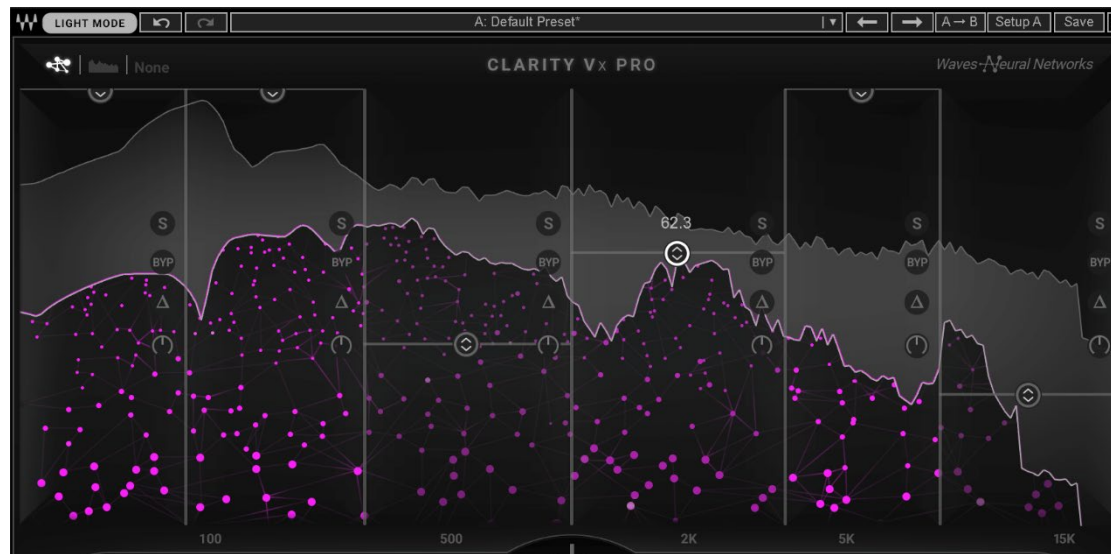
The Link button above the meters locks the faders together. Existing offsets are maintained.



The Show Advanced Controls button controls only the display—settings and controls are not affected. If you have changed any of the Advanced controls from their default positions, and then switched the Advanced button to Off, you will see an asterisk. This indicates that there has been a change in the “hidden” Advanced controls.



Graphic Display



The graphic analyzer display shows the input and output levels of the neural network, as well as the difference between them (reduction). The display is divided into six frequency bands whose ranges are set with draggable crossovers. Each band has a Band Process fader that is used to offset the processing of Main Control knob processing. As you drag a band's process control downward, you reduce the neural network's process strength for that band and bring back some of the ambience.

There are three ways to view the graphic display. Select a view with the buttons at the top of the window.



Neural Network Frequency Analyzer shows input (white outline on top), reduction (grey) and output signals (colorful). The intensity of the background indicates the mix level for each band, ranging from the most colorful (all neural network output) to grayed out. A grayed-out bottom area may indicate a zero setting on the Band Process fader or band bypass. This is the most CPU-heavy view.



Traditional Frequency Analyzer is the Neural Network Analyzer view, but without the sparkly bottom section.



Analyzer Off turns off all analyzer graphing. All band mix controls remain active.



Controlling the process amounts per band

Using the Main Control Knob to set the *Voice vs. Ambience* mix will usually provide excellent noise reduction without impacting the voice. To enable greater control, use the six-band process controls in the Graphic display to adjust processing amounts in defined frequency bands.



- 1 Input signal
- 2 Difference between the input signal and the output signal after processing (i.e., reduction meter)
- 3 Output signal after processing
- 4 Area that illustrates the overall output signal
- 5 Adjustable crossover
- 6 Frequency band
- 7 Band Process fader
- 8 Band solo
- 9 Band bypass (revert to input signal for this band)
- 10 Delta (hear only the reduction for this band)
- 11 Pure Gain control per band



Band Process Controls

CROSSTOVERS

There are three crossovers. Drag a crossover to adjust the boundaries of two adjacent mix bands. It can be helpful to solo a mix band while adjusting the crossover. The crossover frequency is displayed when the crossover bar is clicked or moved.

Default crossover values for six-band presets are 80 Hz, 300 Hz, 1 kHz, 3.5 kHz, and 8 kHz

Default crossover values for four-band presets are 100 Hz, 1 kHz, and 5 kHz

BAND SOLO, BYPASS, AND DELTA

Each band can be soloed or bypassed so that you can better understand how each band is contributing to the mix.

The **Solo** button mutes all other bands. Solo and Bypass buttons are latching, so several bands can be soloed and/or bypassed at the same time. When a band is soloed, the other bands will dim, but their frequency analyzer displays and other controls will still be active.

Bypass bypasses the band processing. This lets you compare the current band mix setting with the input signal.

Delta lets you hear only the input signal, minus output signal (reduction) of the respective band.

PURE GAIN

Use the band **Pure Gain** control to change the gain of a band, post neural network processing. The process amounts do not change when gain is applied, only the level of the band with respect to the other bands. A pink mark is shown on any Pure Gain knob whose value has changed from the default position (0 dB).

Pure Gain between bands has a Q of infinity. This means that there is no slope and no bleed between bands. This enables surgical precision, without the need for smoothing or the fear of clicks or sweeps.

NOTE:

Adjusting the Pure Gain knobs can change the output signal.
These changes are not reflected in the graphic analyzer.



Keyboard Shortcuts and Precision Controls

Keyboard shortcuts enhance workflow and efficiency when working with band controls, enabling further control and precision. We recommend getting acquainted with these shortcuts to get the most out of Clarity Vx Pro.

CONTROL KEY**

Shortcut	Action
Control + Left Click + Drag Crossover	Adjust Crossover frequency while soloing the band to the left of it.
Control + Right Click + Drag Crossover	Adjust Crossover frequency while soloing the band to the right of it.
Control + Click/Drag Gain Knob	Adjust the band's Gain while soloing it.
Control + Click Delta	Monitor the band's Delta while soloing it.

***When using the Control Key shortcut, the "Solo" feature will NOT be recorded to automation.*

Shortcut	Action
Control + Right Click + Delta	Toggles all Delta buttons: first Off, then On.
Control + Right Click + Solo	Disables Solo across all bands .
Control + Right Click + Bypass	Disables Bypass across all bands .



WaveSystem Toolbar

Use the bar at the top of the plugin to save and load presets, compare settings, undo and redo steps, and resize the plugin. To learn more, click the icon at the upper-right corner of the window and open the WaveSystem Guide.

FULL RESET

To reset Clarity Vx Pro to factory default settings, open the WaveSystem presets menu and select “Full Reset.”

LIGHT MODE

The Clarity Vx Pro display is, by default, dark. Click the “Light Mode” button on the left side of the WaveSystem Toolbar to display a light version of the display.

CPU

To avoid running out of CPU in your DAW, don’t open more instances of Clarity Vx Pro than you need. Please also note that Double analysis mode uses substantially more CPU than the Single mode. Therefore, if you are using several instances of Clarity in Single mode and you begin switching them to Double, you may encounter CPU spikes. All of this depends on your host computer, as well as the other plugins that you are running.

The Broad ECO neural network uses significantly less CPU. This provides for far more Clarity Vx Pro instances in a session.

Error Message

If you encounter a message from your DAW or Clarity Vx Pro (such as the message shown below), do not be alarmed. This likely indicates that some essential files required by the plugin were accidentally removed. The solution is usually simple: reinstall the plugin using Waves Central and relaunch Clarity Vx Pro.

