Waves Neural Networks

Clarity Vx DeReverb Pro

Powered by Waves Neural Networks®

User Guide

Contents

Introduction

Thank you for choosing Waves! To get the most out of your new plugin, please take a moment to read this user guide. For even more information, visit this plugin's product page at www.waves.com. We also suggest that you become familiar with the Waves Support pages: www.waves.com/support. There, you'll find technical articles about installation, troubleshooting, system requirements, and Waves contact information.

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 reverb from a dry 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 is 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 the neural network greatly influences its decisions.

Waves Clarity Vx DeReverb Pro

Reverb can enhance the sound of a voice by creating a sense of space, but excessive or inappropriate reverb can create a muddy sound that makes it difficult for the voice to blend well with other tracks. Often, the voice is recorded in sub-par conditions or in untreated spaces, which leads to unwanted reverb and early reflections in the recording. In many situations, there's no choice but to use a vocal track that's already very reverberant, such as an on-location dialogue recording or an a cappella track or vocal sample that comes with baked-in reverb. In such cases the only way to make these recordings usable is to dereverberate them.

DeReverb Pro, from the Clarity Vx series of plugins, enables you to reduce or even eliminate reverb from a voice recording without compromising its voice quality or hurting the ambience and production sound around it. It uses deep neural networks to adaptively isolate a voice from its reverb. DeReverb's extensive controls enable you to quickly and easily adjust the dereverberation curves and control how much reverb to remove or keep. You can select between several neural networks, each trained with different material, to dereverberate speech and sung vocals alike.

In many cases, a single knob is all you need to reach perfect results. You can further precisely control which frequencies to process more and which to process less. Combined, this provides powerful reverb reduction with minimum impact on the voice. For pinpoint control, you can audition the difference between the input and output signals, whether one or more soloed nodes or the entire frequency spectrum. You can control the reverb tail and the presence of the dry output, as well as the stereo width of the output signal. A limiter safeguards against unwanted peaks at the end of the signal chain.

You no longer have to choose between using a track with more reverb than you like or discarding it altogether. You can now reduce or remove unwanted reflections and reverbs, and make impossible recordings come back to life.

Clarity Vx plugins should be inserted at the beginning of a plugins chain. Our neural networks are trained on clean voice tracks, free of excessive compression, EQ cuts, auto-tune, and other processes that make it harder it for the neural network to identify “clean” voice. For this reason, other plugins should not come before Clarity Vx DeReverb.

When using Clarity Vx DeReverb and Clarity Vx together, DeReverb should be first.

Getting Started

This quick start will show you how to separate a voice from its reverb. Follow these steps and you'll achieve a high-quality reverb reduction. Refer to the rest of this user guide to learn more about Clarity Vx DeReverb Pro controls and how to use them for best results.

Processing Sections

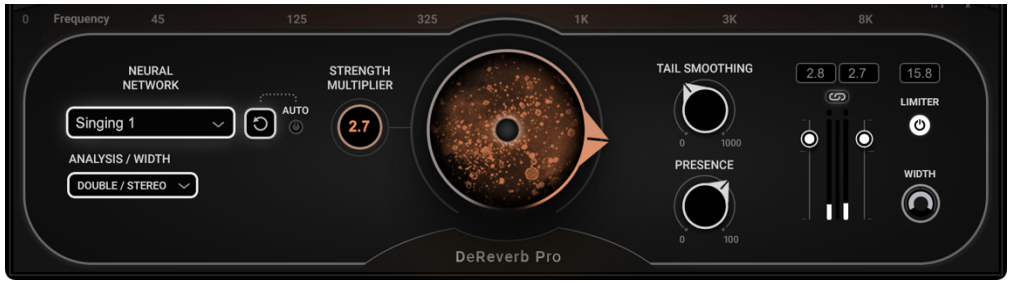

The Main Controls section, located at the bottom of the plugin, is where you configure the settings to remove reverb: select the neural network, set your strength multiplier, increase the tail smoothing, adjust the width of the image, control how much presence you'd like to preserve, and—most importantly—set the amount of neural network reverb reduction.

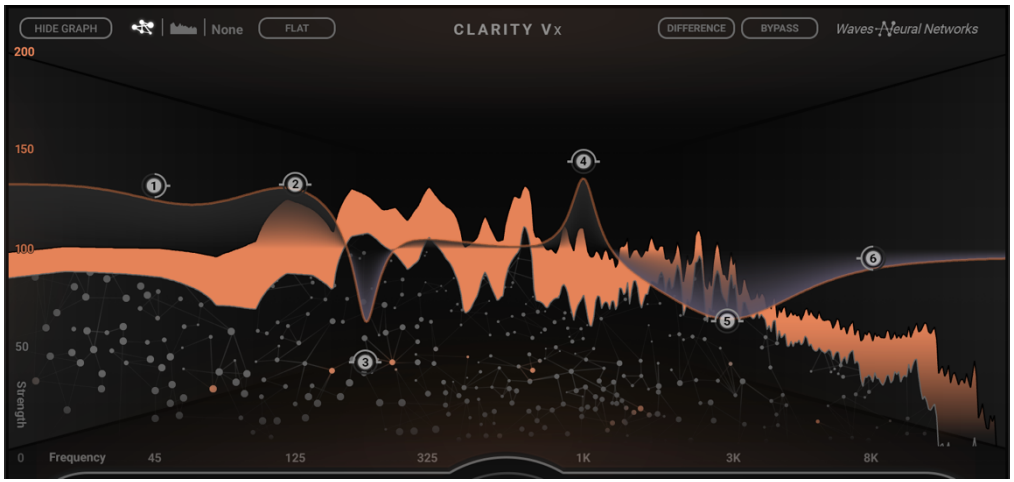

Use the Graphic Controls section at the top to fine-tune dereverberation, in concert with the Main Controls. Adjust the graph for maximum reduction in frequencies with the most reverb and less reduction in frequencies where you want to keep the reverb. This enables strong, focused reduction with minimum impact on the voice.

The Graphic Display can be hidden to save screen space.

Tooltips

Hover over a control for a brief description of its function, which appears in a tooltip on the bottom corner. The value 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.

Interface

Main Controls (always shown)

1. Neural Network Selector, Reset, and Auto Reset
2. Main Knob: adjusts amount of reverb reduction
3. Analysis/Output Width preset modes
4. Strength Multiplier: overall power of NN engine
5. Tail Smoothing: restores some of the original tail
6. Presence control: enhances voice after dereverberation
7. Output section

Graph Controls

8. Input Signal (top line)
9. Output Signal (bottom line)
10. Dereverberated Signal (Diff) (in orange)
11. Dereverberation Strength Curve
12. Node
13. Node Widget Panel
14. Graphic Display view options
15. Difference (Delta) and Bypass

This is a suggestion for getting started. It describes the reverb reduction process and introduces you to most of the controls. If you want to move ahead on your own, just remember two basics: (1) use the bottom panel to establish the texture and amount of reverb reduction, and (2) use the graphic display to shape the process in specific frequencies. The Graphic Display section is strongly affected by the bottom panel.

Quick Start

Step-by-Step: Main Section

- A neural network's training determines what it is best suited for, and Clarity Vx DeReverb Pro has several. Each network has been trained to treat a specific voice type and deal with a variety of background noises that you may want to preserve (e.g., outdoor and indoor ambience, room tone, and other non-voice sounds). Use the **Neural Network drop-down menu** and the associated tooltips to select the most suitable network Try more than one, as results are content-dependent.
- Use the **Main Knob** to set the range of processing. Turn the knob to the right for greater reverb reduction, turn it to the left for less. Experiment until you find the right amount of reverb reduction. If you want to fine-tune the dereverberation processing, use the graphic display at the top. There are six frequency-specific nodes that let you control the power of the reduction. If you plan to use the graphic controls for fine tuning, we suggest that you begin with higher Main Knob values and use the graph to offset processing per frequency range. Using the graph is described in the next section.
- **Analysis/Width** (stereo component only) **Analysis** sets how a stereo signal is analyzed before neural network processing: either as two separate stereo tracks or a mono sum of the two. This setting does not change the audio signal, only the sidechain going into the neural network. **Width**, on the other hand, does affect the output signal and sets the output width between stereo (100%) or mono (0). This can also be done with the Width control in the Output section.

- **Strength Multiplier** increases the power of neural network processing. It is not a mix, like the Main Knob. Rather, it controls how hard the neural network engine is working. It multiplies the processing strength of the nodes in the graph by the indicated multiplier value, which uniformly increases neural network processing in all frequencies. The node settings in the Graph do not change, but the underlying power does. Use this to increase overall processing on particularly difficult reverbs, or experiment with the balance between the Strength Multiplier (neural network "horsepower") and the Main Knob (mix between total Strength and input signal). You can arrive at a very similar reverb reduction, but each with their own texture. The Strength Multiplier is described in the Controls section of this user guide.
- **Presence** adaptively restores closeness and presence of the voice, which may have been lost due to strong dereverberation.
- Use **Tail Smoothing** to instruct the neural network to loosen the release of its process, thus restoring some of the original reverb tail from the input signal.
- Set the output levels using the output **Faders** and apply a **Limiter**. Adjust the width of the output signal with the **Width** control. When working with a mono vocal on a stereo track with a stereo reverb, we highly recommend setting the width to "0."

Step-by-Step: Graph Section

The Graph section allows you to control the amount of neural network processing in specific frequencies. There are six node controls that can be adjusted to any frequency. When you reduce the strength of a node (move it downward), neural network processing is reduced, and there is less reverb reduction at those frequencies. Moving a node upward increases neural network processing, which results in more dereverberation at those frequencies. Graphic controls are subsequent to the Reduction Knob, so adjusting it proportionally will affect total dereverberation. If you plan to use the graphic controls, we suggest that you begin with higher Main Knob and or Strength Multiplier values.

> ! Setting the Main Knob to 100% does not mean that you are at maximum dereverberation. You can achieve greater dereverberation by moving the graph nodes upward, effectively dereverberating the signal even more, even when the Main Knob is set to 100%.

To adjust processing at specific frequencies:

1. Click on a node. This opens a widget that allows you to control the node's parameters.
2. Drag the node to adjust frequency in a freestyle manner. Use the **F** or **Q** controls to lock adjustments to Frequency (horizontal) or Strength (vertical). You can also adjust the Q: solo the node's frequency range or toggle its effect on/off. Nodes are labeled 1–6, but any node can be assigned any frequency. Node 1 is a low-shelf curve and Node 6 is a high-shelf curve. Nodes 2–5 are bell curves.
3. Use the Solo and On/Off buttons to audition the results of the settings.
4. You may want to revisit the Main Knob setting, Tail Smoothing, or presence.
5. Check the output level once again and engage the limiter if the signal goes over 0 dB.

Automation

Every control except the Neural Network Selector, analyzer view options, and collapse/expand can be automated.

> All Clarity Vx DeReverb 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 an internet connection.

Controls

Main Controls

Neural Network Selector

Use the Neural Network menu to choose a neural network that is appropriate for your source material and that suits your project. Each neural network has been trained with different material, and each, depending on the source material, provides different effectiveness in removing reverb, keeping/reducing the ambience, and preserving the voice.

We encourage you to experiment with all the neural networks.

There are many ways to move between neural networks for quick A/B comparisons:

- Click on the neural network name window. Left-click to move to the next neural network, Ctrl+right-click to move to the previous one.
- Use the left/right arrows above the window (these arrows are visible only when you hover over the Neural Network section).
- Use your keyboard's up/down arrows.
- Or, of course, use the drop-down menu.

Main Knob

The Main Knob sets the dereverberation amount. Other adjustments are strongly influenced by it. When the knob is set fully to the left, you hear only the input signal. Turn the knob fully to the right to hear nothing but the neural network processing.

Range: 0% (input signal only) to 100% (neural network output only)

Reset

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 Vx DeReverb resets the neural network every time you stop.

In cases where the reverb type and character have changed between different audio sources in your session, it is best to manually reset the neural network. To avoid a small click or pop, click Reset during an unvoiced region or during silence. This can be automated.

When Auto-reset is on, the neural network automatically resets each time it detects 200 ms of inactivity (levels below -144 dB). This can be beneficial in real-time use when one instance of the plugin is expected to treat multiple events that are separated by at least 200 ms of silence.
Default: off

Analysis/Width

Clarity Vx DeReverb Pro analyzes the input signal before it is separated into voice and reverb. The **Analysis/Width** menu lets you select the analysis method that it is best suited for neural network processing and choose whether the output will be panned fully L/R (100% stereo) or routed to center (mono).

In **Double** modes, the left and right channels are analyzed independently to create a unique set of weights and biases for each channel. In both of the **Single** modes, the left and right channels are summed and analyzed together. Weights and biases provide instructions for neural network processing. The analysis mode affects only the analysis, not the actual signal.

Width sets whether the reverb output signal remains at its original stereo or is routed to mono.

- When Single is selected, both channels are processed identically. This results in a very solid stereo image. If, however, the input signal has significantly different reverbs on the left and right channels, one channel may be underrepresented or overrepresented in the analysis. In such a case, choose Double mode. At any rate, we recommend working with Mono width, when possible.

- Double mode is used when the left and right channels are significantly different (e.g., a stereo reverb, a choir, a group of singers). It analyzes the left and right channels independently and processes them separately. This results in more precise processing for each channel, but it requires more CPU.

In certain conditions when using Double/Stereo, you may find the stereo image to be slightly unstable or “jittery.” In such cases, select Double/Mono or use the Width control in the output section to narrow the plugin’s output signal. This stabilizes the image and produces a cleaner, more solid voice (the stereo component always outputs a stereo signal, so you are creating a mono output on a stereo channel). Mono is always more stable.

When Mono is selected on the drop-down menu, the Width control is disabled. For convenience, you can click repeatedly on the Analysis/Width box to toggle through the menu options.

The Analysis/Width menu and the Width control are available only in the stereo component.

Strength Multiplier

The Strength Multiplier increases the overall power of the neural network, across the entire frequency spectrum. It can be used, for example, when your input signal requires even more dereverberation than the Main Knob and graph nodes can provide. The Strength Multiplier does the same thing as the graph nodes: it increases strength evenly, proportionate to the individual nodes. This gives the Main Knob and the graph nodes much more power for dereverberation.

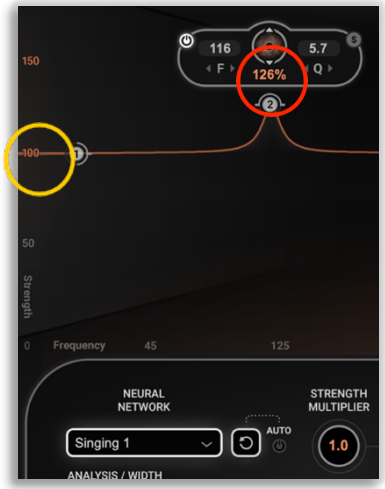

In their default positions, graph nodes are set to 100%, the baseline for the graph (left, circled yellow). With the nodes, you can control the strength range from 0% to 200% (twice the default amount of reverb reduction set with the Main Knob). Note that the strength setting of a node is expressed in percent, not gain (in this case, 126%, circled red).

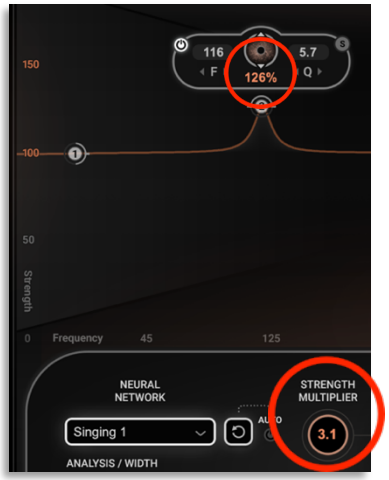

When an individual Node Strength Multiplier is set to a value greater than 100%, the power of the neural network increases, so each node applies more dereverberation processing. In this image, the node value is still set to 126%, but with the Strength Multiplier set to 3.1, there is a significant power increase applied to the graph, and thus greater reverb reduction.

To add strength to specific frequencies (beyond what the notes can provide), increase the Strength Multiplier, which will result in a greater strength on all nodes.

Tail Smoothing and Presence

Clarity Vx DeReverb Pro can completely remove reverb from a track (the early reflections and the reverb tail). However, we are used to hearing a voice with a certain amount of “room,” even in studios. So at times, a dereverberated track may sound too dry or thin, lacking presence in certain frequencies.

Tail Smoothing

Use **Tail Smoothing** to restore some of the original reverb tail. The higher the setting, the longer the tail of the original signal that will be reintroduced to the output.

Range: 0 ms to 1000 ms

Presence

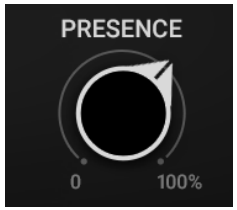

Use the **Presence** control to instruct the neural network how to preserve the voice’s close-mic sound adaptively and automatically. It controls how close and present you would like the voice to be, after dereverberation. A setting of 50 (default) provides a color that closely matches the original. At values of 50 to 100, the neural network gently enhances the sound and adds some high-end (proportionally opposite to its dereverberation process) to further add presence and improve clarity.
Range: 0 to 100

> The Presence setting influences the volume of the plugin. At high values, the output may exceed 0 dB. If this occurs, use the limiter or reduce the output level. Either will control the output gain, but the color will vary somewhat from one to the other.

Output Section

Output faders, post-all processes.

In the stereo component, the left and right output faders can be used independently. Click the Link button above the faders to unjoin them and adjust individually. Click Link again to rejoin the faders, while maintaining any offset between them.

Range: -inf to +24 dB

Link default: On

Meters show output signal, post all processes, with red peak-hold LED**.** Click on the meter to reset the clip LED.

Limiter

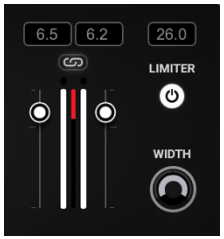

A brickwall limiter can be applied at the end of the signal chain to protect from harsh clipping artifacts and ensure the set output level is not exceeded. When the limiter is active, a Gain Reduction meter and Gain Reduction peak hold value are displayed.

Range: On or Off

Default: Off

Width (stereo component only)

Adjusts the size of the stereo image.

Range: 0% (mono) to 100% (original stereo width)

Default: 100%

Graphic Display

The Graphic Display enables fine-tuning of the overall reverb reduction that has been defined with the Main Knob. Reverb reduction can be increased to further reduce particularly challenging reverb in specific frequencies without affecting other frequencies. Reverb reduction can be lessened in specific frequencies to keep some of the original reverb. This provides the greatest possible reverb reduction and control, with the least possible artifacts.

The display shows the input and output levels of the neural network, as well as the difference between them (reverb reduction shown in orange). The graph includes six nodes, each of which can control frequency, strength, and Q.

As you drag a node downward, you reduce the neural network's process strength for those frequencies and bring back some of the reverb.

When you move the node above the default center line, dereverberation increases at the specified frequency. This can result in very aggressive processing around that node, and a very dry signal, so you may want to compensate with the main knob.

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 (top line), reduction (orange) and output signals (bottom line, on top of the sparkles). This view may require slightly more CPU.

Traditional Frequency Analyzer is a more streamlined analyzer view, without the sparkly bottom section.

None turns off the analyzer display. All other graph controls remain active and will not change while the graph is hidden.

Show/Hide Graph

You can hide the Graphic Display section to save space on the display. Click the **Hide Graph/Show Graph** button in the upper-left corner of the window.

By default, Clarity Vx DeReverb Pro loads with the graph visible, but it will remember your last view on the next instantiation of the plugin. If any Graph settings have been changed while the graph was hidden, an asterisk will appear next to the Show Graph button to indicate that the graph controls are not at their default state.

Controlling Strength by Frequency

Using the Main Knob to remove reverb will usually provide appropriate reverb reduction without impacting the voice. To enable greater control, use the graphic controls to adjust processing amounts per frequency ranges.

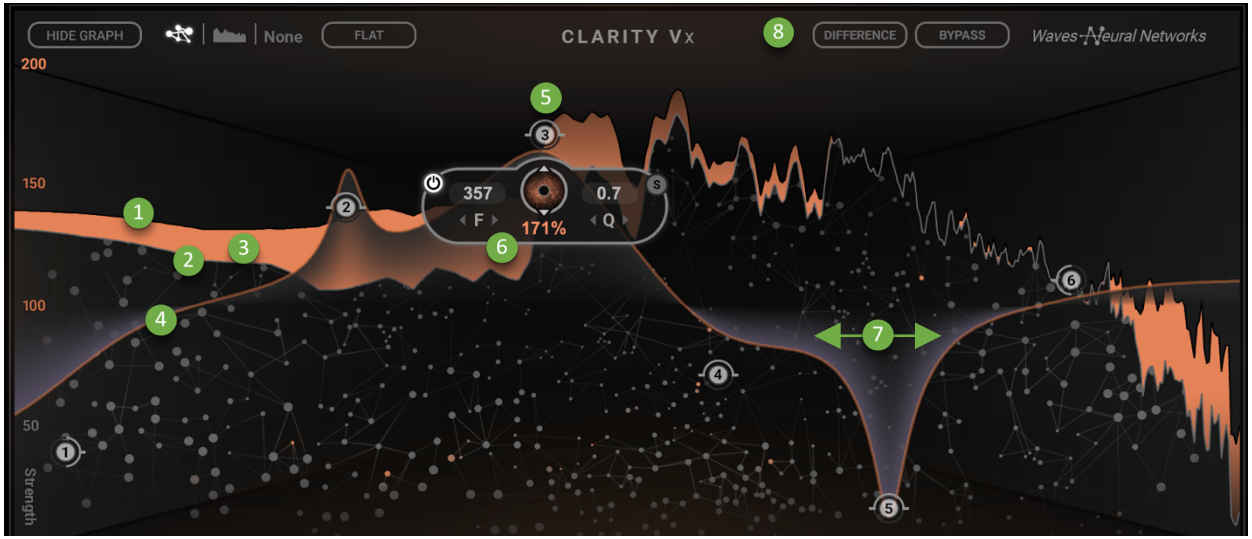

1. Input signal
2. Output signal after processing
3. Difference between the input signal and the output signal after processing (reverb reduction/difference)
4. Outline (curve) of the total dereverberation strength settings
5. Node: drag for freestyle parameter adjustment, click to open its widget
6. Node widget: controls node parameters
7. Q: range of processing per node
8. Bypass and Difference on/off

Adjusting the Nodes

The shape of the graph curve sculpts the reverb processing by user-defined frequency regions. The graph is drawn with six nodes that govern the frequency, strength, and Q. Any node can operate between 16 Hz and 21,357 Hz. These node parameters determine how the deverberation will be fine tuned by frequency.

Node curve types are fixed:

Node 1	low shelf
Nodes 2–5	bell curves
Node 6	high shelf

Node Widgets

A node is controlled from a widget. Click on a node marker to open the widget.

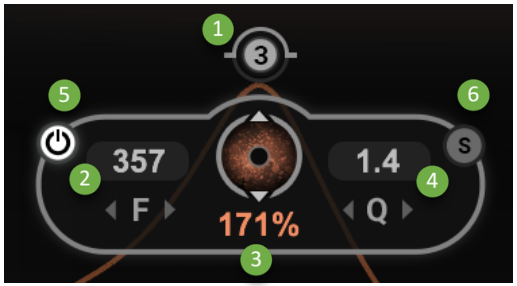

1. Node Marker (click to open widget; drag marker for freestyle control)
2. Frequency adjust and value
3. Strength adjust and value
4. Q adjust and value
5. Node on/off
6. Node solo

Using the widget

Keep in mind that the greater the strength, the greater the reverb reduction at that frequency, and the less reverb you will hear. Lower strength settings mean less reverb redcution, and more of the original reverb will be kept.

- Drag a node **marker** (not the widget) to adjust frequency and strength simultaneously, without constraint.
- Drag the **F** control on the widget to restrict the adjustment to frequency only (x-axis).
  Range: 16 Hz to 21,357 Hz
- Drag the **Strength** button to restrict the adjustment to Strength only (y-axis).
  Range: 0% to 200% (default 100%)
- Drag the **Q** control to restrict adjustment to Q.
  Range: 0.4 to 7 (default 0.7)
- Click on **S** to solo the node; click **On/Off** to turn off processing for that node.

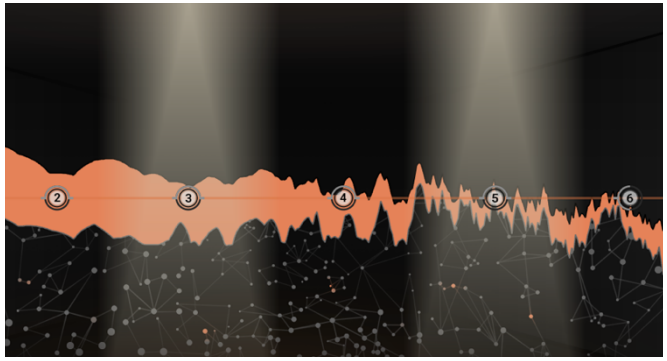

Solo auditions only the frequency range of the node. All frequency ranges outside the soloed range (or ranges) are muted (-inf dB). A flashing node indicates solo, and a spotlight illustrates the soloed frequencies. In this example, Nodes 3 and 5 are soloed.

On/Off turns off dereverberation processing in the selected node, which reverts to its default values as long as it is set to off. Once turned back on, it will recall its previous settings.

Difference and Bypass (Global)

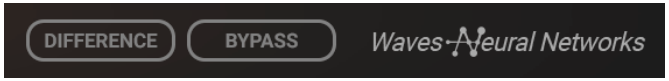

These controls (at the top of the Graphic section) apply to the entire plugin.

Bypass has two behaviors:

- *When no nodes are soloed*, click Bypass to hear only the input signal. This is the same as using the DAW's bypass function.
- When one or more nodes are soloed, click Bypass to hear only the input signal of the *soloed notes*. This helps you understand how the processing affects the input signal at specific frequencies.

Difference auditions the difference between the input and output signals, essentially playing only the reduced reverb. *When one or more nodes are soloed*, Difference auditions the difference between the input and the soloed node(s).

Flat

Reverts all nodes (the entire graph) to their default settings and locations without changing other controls.

Special uses for nodes 1 and 6

Use the shelves in Node 1 and Node 6 to change the focus and amount of dereverberation, without adjusting individual nodes.

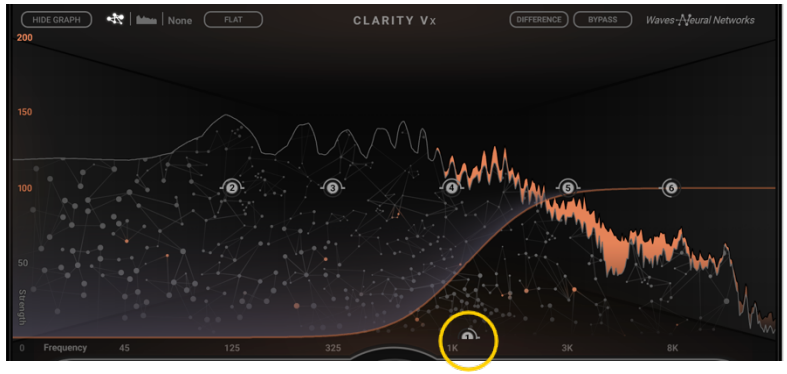

In this example, Node 1 (circled yellow) is moved to the right and down to eliminate almost all dereverberation. Since this node is a low-shelf curve type, dereverberation is reduced by frequency as the node moves right and by strength as the node moves down.

Here, Node 6, a high-shelf filter, is moved to the upper-left corner. This increases the strength of all nodes, without changing their individual settings. As in the previous example, you can move Node 6 sideways to control its behavior.

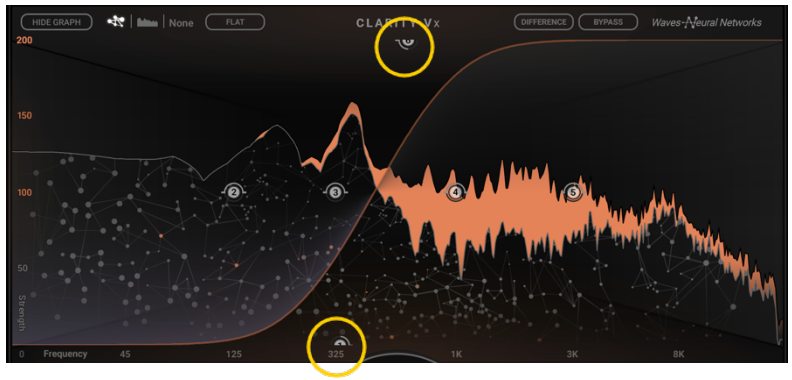

You can use Nodes 1 and 6 together to achieve maximum dereverberation of the entire high-end while preserving the low-end reverb. Drag Node 6 up and to the left and Node 1 down and to the right. This does not change the settings of Nodes 2–5. This shelving action “sits on top” of the other nodes.

Control Shortcuts

Node Controls

Adjust Strength (Freq locked)	Ctrl+drag node up/down
Adjust Strength (Strength locked)	Ctrl+drag node left/right
Adjust Q	Shift+drag *or* mouse wheel/track pad two fingers up or down on a node
Node on/off (value reverts to 1.0)	Double-click on the node
Solo a node	Ctrl+right-click on the node

All Solos off	Ctrl+right-click on any Solo

Multi-Select

All controls can be multi-selected for grouped control, except for widget controls, neural network selector, and reset button.