CEDAR Audio Limited

www.cedaraudio.com

This page left blank.

Table Of Contents

This page left blank.

Getting Started

CEDAR Icons™ are award-winning processes derived from groundbreaking research and our flagship CEDAR Cambridge™ system. These can process your audio in many desirable ways and are capable of enhancing your sound and correcting all manner of problems without damaging the desired audio or introducing unwanted side-effects and artefacts.

Unpacking and licensing

CEDAR Icons are supplied electronically for Mac and PC hosts, and authorised using an iLok or iLok Cloud, so no physical objects are necessary or supplied. Licensing is performed using the iLok licence manager. Please refer to your iLok documentation.

Assumed Knowledge

CEDAR Icons conform to the 64-bit AAX Native and 64-bit VST 3 and AU plug-in formats. They are not compatible with 32-bit operating systems or hosts.

This manual assumes that you are fully conversant with your host computer, and that you know how to operate the host software into which you're loading your Icons. It may refer to operations that are common to these products, but will not attempt to explain them.

Troubleshooting Non-CEDAR Components

If you encounter problems with your Macintosh®, macOS®, your PC, Microsoft Windows®, or any 3rd party host system, please refer to the relevant manuals or contact the dealer that supplied these to you. Unless appointed independently as authorised dealers for these, CEDAR Audio's dealers will not attempt to provide technical support for them.

Installation

Mac:

- Before proceeding, ensure that any older versions of the plug-ins or CEDAR Studio versions (if any) are removed from their plug-in folders.
- If you are unable to locate the uninstallers, you may remove older plug-ins manually from the following folders found in: **Applications/CEDAR Audio Ltd**

VST 2 plugin files:	HD/Library/Audio/Plug-Ins/VST
VST 3 plugin files:	HD/Library/Audio/Plug-Ins/VST3
AU plugins files:	HD/Library/Audio/Plug-Ins/Components
Pro Tools AAX plugin files:	HD/Library/Application Support/Avid/Audio/Plug-Ins

- Once you have removed and older versions, double-click on the installer:

  CEDARAudioLtd-xxxxx.pkg

- Follow the instructions.
- If desired, you can click on the Customise button to select the plug-in format or formats (AAX, AU, VST3) that you wish to install.
- Ensure that you have sufficient space on your drive and then press Install to install the complete package or press the Customise button to select which formats are loaded.
- A message will appear to tell you that the installation was completed successfully.

PC:

- Before proceeding, ensure that any older versions of the plug-ins or CEDAR Studio versions (if any) are removed from their plug-in folders.
- If you are unable to locate the uninstallers, you may remove older plug-ins manually from the following folders found in: **C:\Program Files/CEDARAudioLtd**

VST 2 plugin files:	*C:\Program Files\Steinberg\VSTPlugins*
and/or	*C:\Program Files\Common Files\VST2*
VST 3 plugin files	*C:\Program Files\Common Files\VST3*
Pro Tools AAX plugin files:	*C:\Program Files\Common Files\Avid\Audio\Plug-Ins*

- Once you have removed and older versions, double-click on the installer:

  setup.exe

- Follow the instructions.
- If desired, you can click on the Customise button to select the plug-in format or formats (AAX, AU, VST3) that you wish to install.
- Ensure that you have sufficient space on your drive and then select the plug-in formats you wish to install.
- A message will appear to tell you that the installation and, if selected, the driver update was completed successfully. Click on OK and Finish.

General information

CEDAR Quantum™ near-zero latency technology

Latency is the most significant barrier to using plug-ins for live sound, so many CEDAR Icons have been developed with near-zero latency that allows you to use them in live situations. If you see the CEDAR Quantum logo on the interface, the internal latency of the plug-in is less than one millisecond, so you can insert it in any audio stream without adding any perceptible delay in a concert, theatre or auditorium, and without fear of losing lip-sync when broadcasting.

Bypass

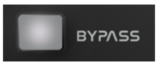

All CEDAR Icons feature a Bypass button. This directs the input signal to the output. The process is still active in the background so you can remove the bypass without generating a significant discontinuity.

Automation

General

Icons' controls appear in the plug-in automation screens of appropriate host software and may be automated in standard fashion.

DNS One

When Learn is On and your host system is recording automation, the Level and Gain values calculated within DNS One are provided to the automation system and recorded if the host is capable of doing so.

When Learn is On and the host system is replaying automation, the Level and Gain values calculated within DNS One are used and those replayed by the automation system are ignored.

Backward compatibility

Where appropriate, VST and AAX Icons are backward compatible, so you can reload sessions that used the equivalent CEDAR Studio processes.

To preserve compatibility with GarageBand, Icons are not backward compatible with sessions that used AU versions of CEDAR Studio.

Knobs

Double-click on any knob to return it to its default value.

Readouts

Double-click on any numeric readout to edit its value directly.

Adaptive Limiter 2

Everything from complete transparency to heavy maximising

Dynamics processes are used extensively in a creative fashion to obtain desirable characteristics within a given piece of audio. Most commonly, they are used to increase the average signal level (with respect to the maximum signal level) and to set that maximum to a desired level - often 0dB or a little below. Capable of doing both simultaneously, CEDAR's Adaptive Limiter 2 fulfils the functions of flexible and transparent peak limiter to a powerful loudness maximiser.

Why "Adaptive"?

A conventional limiter controls the loudness of a signal by observing the peak amplitude and suppressing this when it would otherwise exceed a user-defined threshold. A multi-band limiter does the same thing in discrete parts of the spectrum after dividing the signal into multiple bands. Adaptive Limiter 2 does neither of these, instead using a unique algorithm that calculates a continuingly varying EQ profile that constrains the amplitude of the output while retaining the integrity of the input signal.

Controlling the attenuation profile

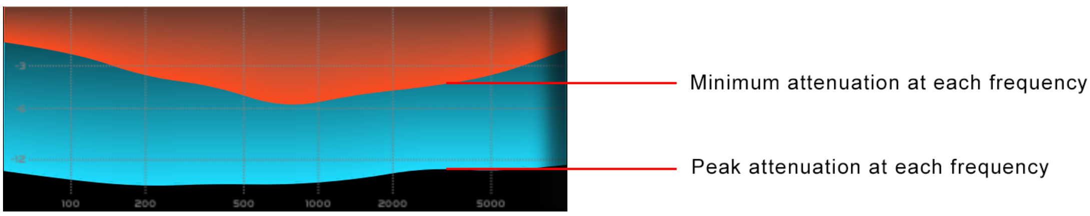

The attenuation profile is constrained in two ways: the rate at which it is permitted to vary in time, and the amount by which it is permitted to vary from frequency to frequency across the spectrum. The three Adaption parameters control this.

The action of the Adaptive Limiter is shown by the two lines in the display.

- The upper line (the boundary between the orange and the blue regions) shows the minimum attenuation at each frequency within a small temporal window.
- The lower line (the boundary between the blue and the black regions) shows the peak attenuation at each frequency within a small temporal window.

Spectral

As you increase the Spectral value, you are allowing the attenuation profile to change more from frequency to frequency. As a consequence, the lines may develop more bumps as the value increases.

Temporal controls

The two temporal controls constrain the rate at which the attenuation profile changes at very low frequencies (LF) and very high frequencies (HF). If the two are not the same, the algorithm generates a suitable profile across the spectrum. The distance between the lines in the display will increase as the Temporal values increase. As you increase the LF temporal value, the process allows you to obtain greater perceived loudness, but at the possible expense of (possibly pleasing) distortion.

This following table explains the responses obtained at the upper and lower extremes of these controls. A wide variety of useful responses may be obtained between these extremes.

	Spectral = low	**Spectral = high**
Temporal LF & HF = low	**The response is maximally constrained in both time and frequency.** The attenuation will be close to flat across all frequencies and will vary relatively slowly. The difference between the peak attenuation and the minimum attenuation at any given frequency will be small. This is similar to the action of a single-band limiter with a slow release. 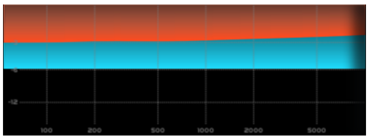	**The response is minimally constrained in frequency but maximally constrained in time.** The attenuation is allowed to vary considerably across the spectrum, but will vary relatively slowly in time. The difference between the peak and minimum attenuations at any given frequency will be small. This is similar to the action of a multi-band limiter with a slow release. 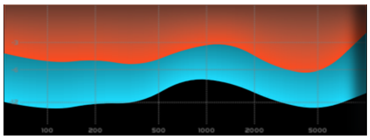
Temporal LF & HF = high	**The response is maximally constrained in frequency but minimally constrained in time.** The peak attenuation will be close to flat across much of the spectrum but may vary quickly. The difference between the peak attenuation and the minimum attenuation at any given frequency will be greater than if the temporal response were slower. This response is similar to the action of a limiter with a frequency-dependent response. 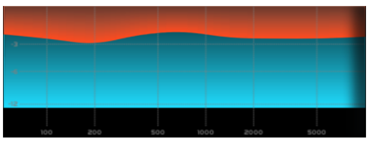	**The response is minimally constrained in both time and frequency.** The attenuation profile will vary considerably across the spectrum and will be free to vary quickly in time. The difference between the peak attenuation and the minimum attenuation at any given frequency will be greater than if the temporal response were slower. These settings maximise the perceived loudness of the output with an increased change in the character of the output. 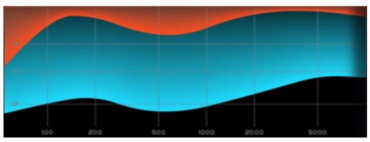

Controlling the gain

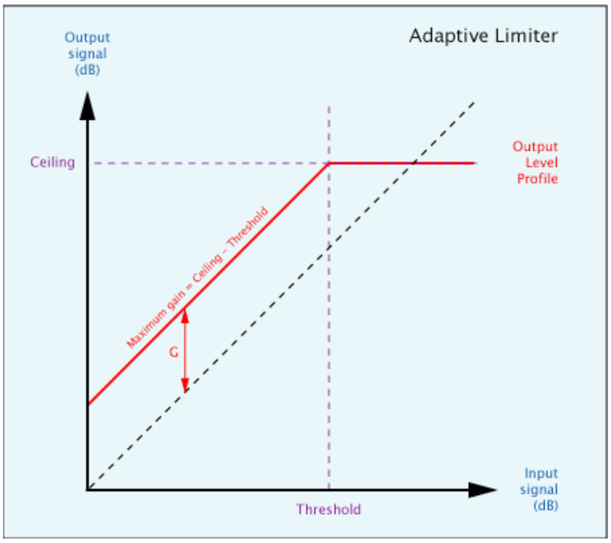

This illustrates the relationship between the input signal level and the output signal level, and how this is affected by the Threshold and Ceiling controls. These can be defined as follows:

Threshold

The input signal level above which limiting occurs.

Ceiling

The maximum output signal level, which is the output level of any limited signal.

Oversampling

With Oversample ON, the incoming signal is oversampled and all subsequent calculations are carried out taking into account the inter-sample peaks.

Quantisation

Adaptive Limiter 2 can quantise and noise-shape its output at a desired word length. These allow you to use it at the end of your processing chain as a high-quality mastering limiter.

Select the following options from the drop-down menu:

Word length and dithering

Dithering and noise shaping are applied whenever OFF is **not** selected, even if the selected word length is the same as that of the input audio. You may select the output word length to be either 16-bit or 24-bit.

Noise shaping

Select one of the four options for your word chosen length. Choose 'flat' for TPDF dithering, or one of the three other forms of noise shaping as desired.

No re-quantisation applied

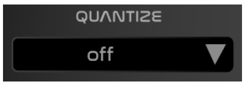

When OFF is selected, no re-quantisation is applied.

Tutorial

The peak level of a digital waveform determines the maximum amount by which you can amplify it before clipping occurs. However, many signals contain high peak levels of short duration, and these can restrict the RMS (average) level to a low figure. The ratio of peak level to RMS level is called the 'crest factor' of the waveform. Since, to a first approximation, it is the RMS level and not the peak level that determines how loud we perceive a sound to be, signals with high crest factors can sound disproportionately quiet. To overcome this, limiting can be used to reduce the crest factor and thus increase the apparent loudness of the signal without increasing the peak level.

Limiting a signal in this way can introduce unwanted distortion, but Adaptive Limiter 2 is able to constrain the peak level and increase the apparent loudness in a wholly transparent fashion. When pushed harder, it generates euphonic distortions that many users find desirable. A limiter capable of doing this is often called a loudness maximiser.

Processing audio

Pass some audio through Adaptive Limiter 2.

Set the Ceiling to 0dB by adjusting the knob or sliding the marker in the Output axis of the metering. It is common to set the Ceiling (or Limit) of limiters to a little below 0dB to ensure that no clipping occurs but, when oversampling is switched on, you can safely set Adaptive Limiter 2 to 0dB if that is the maximum level desired.

Now set the Threshold by adjusting the knob or sliding the marker in the Input axis of the metering such that the level of the input signal is sometimes (but not always) higher than the Threshold. When this occurs, you will see the curve appear in the attenuation profile window, and limiting is being applied.

Decrease the Threshold. You will see increased activity in the attenuation profile window and the loudness of the signal will increase. This is because the difference between the Ceiling (in dBs) and the Threshold (in dBs) is the gain applied to any input signal lying below the Threshold. Any input signal above the Threshold will be limited.

To obtain maximising, decrease the Threshold still further. You will start to hear a thickening of the signal accompanied by mild distortion. Pushed still further, you can generate creative distortion effects.

Controls

<table>
<tr><td>


</td><td>

Level Controls

Threshold

Determines the input amplitude at which limiting is applied.

Ceiling

This sets the maximum level output by the limiter.

Note: You can also adjust these parameters by sliding the markers in the axes of the metering display.
</td></tr>
<tr><td>


</td><td>

Spectral Adaption

Spectral

Determines the degree by which the profile can change from one frequency to the next. At its lowest setting, the action of the Adaptive Limiter approaches that of a single-band limiter. At its highest setting, it is free to calculate a complex attenuation profile.

Temporal (LF and HF) controls

Determines the rate at which the profile is permitted to change at very low frequencies (LF) and very high frequencies (HF). If the two are not the same value, a suitable profile is generated across the spectrum, allowing you to obtain a wide range of limiting effects.
</td></tr>
<tr><td>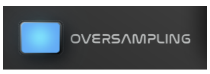
</td><td>

Oversampling

Turn this on to take into account inter-sample peaks in the incoming signal.
</td></tr>
<tr><td>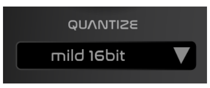
</td><td>

Quantisation

Select the desired word length and noise shaping for the output signal.
</td></tr>
<tr><td>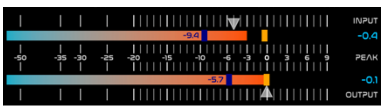
</td><td>

Metering

View the input and output levels as well as the values for the Threshold and Ceiling. You can adjust the values of the Threshold and Ceiling by sliding their markers left and right in this display.
</td></tr>
</table>

Declick

Banishing clicks and ticks

Declick removes clicks from a wide range of material without introducing unwanted distortion or artefacts. For each piece of material that you process, you should determine the best value for the threshold.

Controls

	**Threshold** The threshold ranges from 3 to 100 on an arbitrary scale and controls the sensitivity of the process. With the threshold set high, Declick will remove only the largest clicks and scratches. A lower threshold will also remove smaller ticks and clicks. You should attempt to find the highest value at which the unwanted noise is removed.

Declip 2

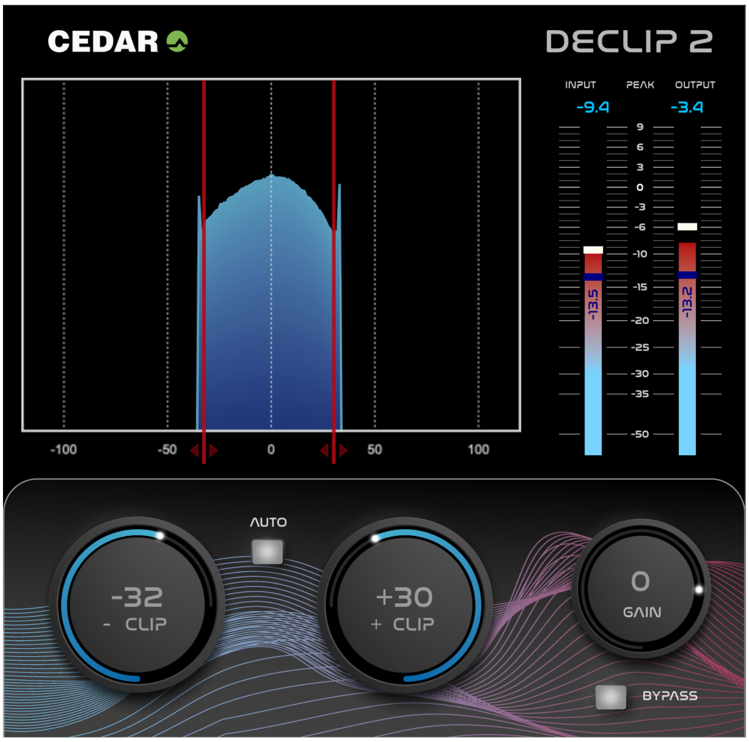

Removing clipping distortion to recover the original signal

Clipping occurs when a piece of equipment or a medium carrying a signal is unable to handle a high level presented at its input, and will usually be heard as a harsh distortion that increases in intensity as the clipping becomes heavier. In all cases of clipping, it is the portions of the waveform near its extremities that are affected, while portions of the waveform closer to zero are unaffected.

Viewing the audio on an amplitude vs. time graph shows flat tops and bottoms where the genuine signal has been destroyed and replaced by false samples at the maximum amplitude. Attenuating the signal after clipping has occurred will affect the amplitude (and, therefore, the sample value of the flattened regions) but will not restore the original data.

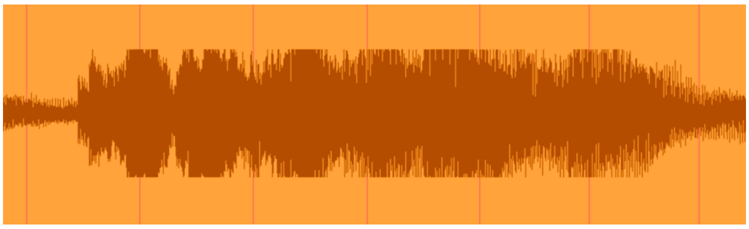

Declip 2 removes the false samples and replaces them with samples that represent the signal that would have been recorded had clipping not occurred.

Declip 2 Tutorial

Using the Signal Analysis Window

If you observe the sample density curve of unclipped audio you will see that it exhibits a reasonably smooth distribution across the sample range. An example of this appears below.

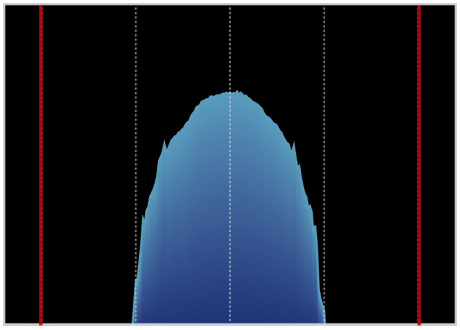

When the audio is clipped, this curve assumes a truncated shape, with peaks at one or both edges. This is because the number of times that this sample value has occurred is unexpectedly high. In the case of hard clipping these peaks will be sharp and well-defined; in the case of soft clipping they may be more rounded, so analogue clipping will usually be less obvious than digital clipping. Nonetheless, the method for identifying and correcting the problem remains the same.

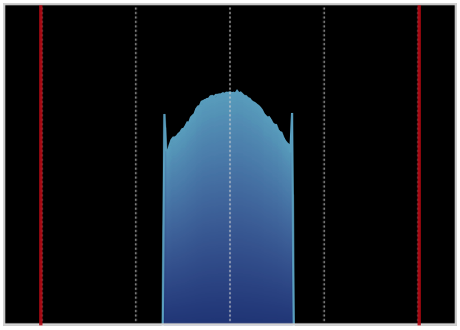

Removing the clipping

When you are happy that you can see the sample values at which clipping has occurred, you should use the **- Clip** and **+ Clip** knobs to move the Clip Markers 'inside' that amplitude. In many cases, you will see that there are valleys that differentiate the damaged data from the good data, and these will most likely be the ideal places to position the markers. Clipping is not necessarily symmetrical, so the - Clip and + Clip values do not need to be the same.

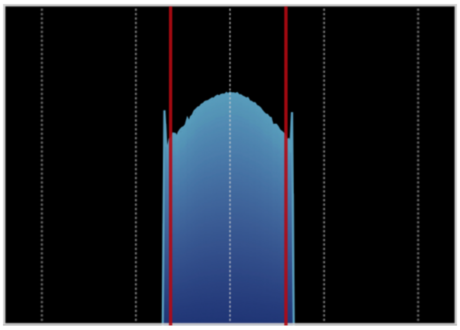

Alternatively, the Auto function will allow the algorithm to determine these values and to track changes in the clipping level, which may be of benefit if any dynamics processing has been applied to the audio after clipping occurred.

Gain

The peak amplitude of a declipped signal can be significantly greater than that of the incoming signal. To output this without reintroducing clipping, it may be necessary to attenuate the restored signal. Use the Gain control to fit this into the available headroom.

The associated meters show the peak amplitudes of the incoming (lefthand) and outgoing (righthand) signals.

Controls

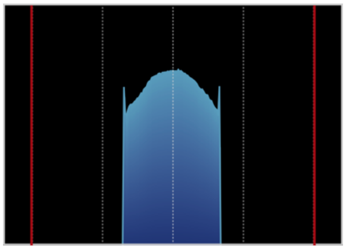	**Sample density curve** This displays the signal represented as the likelihood of a given sample value (the vertical axis) plotted against the sample value (the horizontal axis). Sample values that occur frequently will have greater values on the vertical axis than those that appear on few or no occasions. Clipping is visible as peaks at one or both sides of the display.
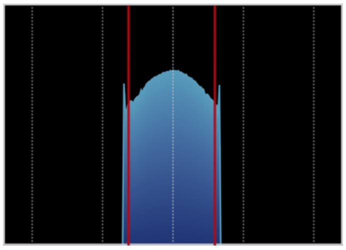	**Clip markers** The Clip Markers are vertical red lines. Any samples lying outside these (i.e. further from the centre line) are removed and replaced with restored signal.
	**Clip detection controls** Use these to move the Clip Markers into position. **Auto** Click on the Auto button to allow Declip 2 to place the markers and track any changes in the clipping level.
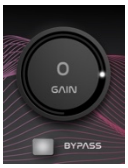	**Gain** Allows you to attenuate the restored audio to fit the available headroom.
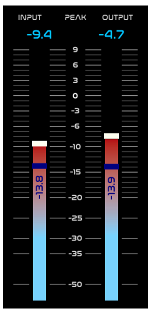	**Meters** Show the instantaneous and peak levels of the input signal (left) and the restored signal (right).

Declip 2 is a further development based upon the Sparse Audio Declipper algorithms developed by Inria - Panama research team.

Decrackle

Banishing crackle and similar unwanted problems

Decrackle is based on a process invented by CEDAR called 'Split & Recombine'. It removes surface noise, crackles, and some types of buzz. You may also use it to reduce the audible effects of some amplitude distortions.

It will only function correctly if there are no major ticks or clicks in the audio. You must remove these using CEDAR Declick before processing the audio using Decrackle.

The threshold controls the sensitivity of the process. A high threshold tells the system to remove only the most obvious crackles and buzzes, while a lower threshold also removes fine crackle, buzz, and distortion.

Controls

	**Threshold** Find a threshold level that removes as much of the crackle/distortion as possible without introducing unwanted side-effects. Low thresholds (4 or 5) can be most effective, but a higher threshold is advisable if processing begins to introduce side-effects. In general, no artefacts will be introduced with a threshold of 8 or above.

DNS One

Near-zero latency dialogue noise suppression for post

The standard for dialogue noise suppression in film and TV post, DNS One offers the Academy Award® winning processing of the DNS1000 and its successors, but adds CEDAR's unique LEARN capability that allows the processes to adapt to changes in the noise, even when the wanted signal is present. It makes unusable interviews intelligible, saves huge costs in ADR, and has rescued dialogue for countless movies. It uses CEDAR's zero-latency Quantum technology so that there is no loss of lip-sync when streaming audio through the process.

Range Selectors

The Range Selectors concentrate the unit's activity into the desired part of the audio spectrum. Selecting any of these concentrates all of the filters within DNS One's filter bank across that part of the audio spectrum.

Level Control

The Level Control tells DNS One how much noise is present in the input signal.

Band Gain Controls, Numerical Readouts and Bargraphs

DNS One divides the chosen range into a large number of well-defined bands. Sophisticated digital filters analyse each of these and suppress the noise independently in each. The innovative design of this filter bank allows you to adjust DNS One using just six controls.

Band gain controls

These determine the maximum amount of noise attenuation that DNS One can apply in each band.

Numerical readouts

These show the values for each fader.

Bargraphs

These offer a visual indication of the activity in each of the bands.

Learn

When Learn is On, DNS One calculates an estimate of the Level and determines suitable Band Gains. You will then be unable to adjust these, although you may still choose the Range in which processing is applied.

Learn is not a noise fingerprint. You may switch it on/off at any time and it will adapt to the noise in the recording. You do not need to find a section of the audio that contains little or no wanted signal. Nevertheless, you can use Learn to take a snapshot of the Level and Band Gains and then fine-tune these manually if you wish.

Reset buttons

These allow you to set all six band gain controls to 0dB (no noise attenuation at any frequency within the range) or to -24dB (maximum noise attenuation at all frequencies within the range).

Tutorial

The following tutorial and case studies teach you how to adjust DNS One's settings manually.

Range Selector

Your first job for any noise reduction will be to identify the frequencies in which the unwanted noise lies, and then select one of the six possible ranges. With practice, you will be able to do this by ear. Until then, you may prefer to use the methods described in the following case studies.

Level Control

Next, you should identify the noise level within the audio. With the appropriate range selected, use the lower button to pull all six Band Gain controls down fully to -24dB. Starting with the Level fader at -80, increase it slowly. At first, you will hear very little happen but, at some point determined by the noise content of the recording, you will hear the noise disappear. You should attempt to determine the point at which this occurs. (When the Band Gain controls are fully down, maximum processing occurs for any given Level setting, thus making identification of the noise easier.)

Band Gain Controls

Once you have chosen the Range and determined the Level, you control the action of DNS One's filter bank using the Band Gain controls.

The six faders represent six frequency bands (each containing multiple filters) distributed from lower frequencies (left) to higher frequencies (right) across the selected range.

You use the Band Gain controls to control the amount of noise attenuation performed in each band, adjusting them to suppress as much noise as possible without introducing unwanted artefacts into the desired signal.

Noise suppression occurs in a given band when the fader is below 0dB. However, there may be occasions when you wish to boost the noise in a given band, and you can do this by moving the appropriate fader above 0dB. This is a particularly useful way to find noise that's concentrated in a narrow band of frequencies.

You can move all the Band Gain controls to 0dB or to -24dB by pressing the appropriate reset button.

Case studies

The following three examples illustrate ways to use DNS One. They may not be the ways that you choose to operate it for all jobs, but they will get you started.

Suppressing traffic noise and other ambient sounds

DNS One can suppress sounds such as road traffic, aircraft, air conditioning, wind and rain, plus many other common soundstage, location, and OB (remote) noise problems.

First, identify the frequency range(s) in which the noise lies

You should be able to do this by listening to the problem. If this proves difficult, you can use the following method. It is not important that you find the perfect settings at the first attempt. In particular, you will be able to refine your Level and Band Gain settings once you have found the correct range.

Start by ensuring that DNS One is not in Bypass and then select Full range processing. Now set all six Band Gain controls to -24dB by clicking on the lower button.

Next, move the Level control to -80 and then raise it until the noise disappears. At this point you have determined an approximate setting for the Level. This is necessary for determining the range but it is likely that you will refine this later in the procedure.

When the Level is close to the ideal setting, you should see the bar graphs flicker in response to the signal content. Now raise the Band Gain controls to 0dB by pressing the upper button (no processing occurs) and adjust the Band Gain controls individually to suppress the noise. You should always attempt to suppress the noise with the minimum of damage to the desired signal.

In many cases, you will find that the leftmost Band Gain controls are pulled down significantly, whereas the central and rightmost are closer to 0dB. This tells you that the problem does not lie in the upper frequencies so you should now select Low+Mid and repeat the previous steps.

If you now find that you are using all six faders, it's likely that the noise is distributed across the Low+Mid range. However, if the suppression is still heavily biased towards the lefthand faders, you can select the Low range and repeat the procedure.

Second, optimise the Level control

Return the Band Gain controls to -24dB and, listening carefully to the audio, refine the Level setting so that the noise is identified correctly. Note that DNS One takes a short period to settle after moving the Level control (especially in the lower ranges) so you should not adjust it too rapidly.

Third, refine the Band Gain controls for optimum suppression

Press to reset all six Band Gain controls to 0dB. Now increase and decrease the Gain in each band separately while listening to the effect that this has on the noise. This will identify the bands that contain the majority of the noise. Don't be alarmed if all six bands contain significant noise. This is not unusual.

Let's assume that the greatest improvement occurs when you reduce the gain in bands 3 and 4. This suggests that the noise is concentrated in an approximate range of 200Hz to 1kHz. You should now find the optimum positions for all six faders. The greatest cuts will lie in bands 3 and 4, whereas bands 1, 2, 5 and 6 should remain as close to 0dB as possible to ensure that minimal processing occurs in the other bands. The final configuration might look something like this:

Suppressing tape hiss

The DNS One can suppress the tape hiss that mars many older recordings. It will also improve the signal/noise ratio of dialogue tapes that have been poorly copied as well as those that are many generations old.

First, identify the frequency range(s) in which the noise lies

You should follow the procedure laid down in case 1 (traffic noise) to determine the range in which the problem lies. For most instances of tape hiss, you will find that the Mid+High range is most appropriate. In a few cases you may find that the High range alone is more suitable. It is not as common to require suppression in the Low range because hiss is usually less prominent at lower frequencies, and it may also be masked by the wanted audio in the range.

Second, optimise the Level control

To determine the correct Level, you should follow the procedure described in the first tutorial.

Third, refine the Band Gain controls for optimum suppression

As in the first tutorial, you should start with all six Band Gain controls at 0dB and then increase and decrease each control individually to find the bands that contribute most hiss to the signal. Because tape hiss often exhibits a flat profile at Mid and High frequencies, you may find that satisfactory results are achieved with the Band Gain controls set in a horizontal line.

However, hiss is generally less annoying at very high audio frequencies so you may be able to reduce the amount of processing in the uppermost bands. This will help to ensure that any low amplitude signal components lying at these frequencies (which provide much of the 'air', 'ambience' or 'life' in a signal) are passed with little or no attenuation.

The final configuration might look like this:

Suppressing excessive reverberation

In many situations, DNS One can suppress excessive reverberation. This can be useful in TV production when you need to match the audio from a large recording studio or soundstage to visual images set in a small room or other enclosed space. Suppressing reverberation can also be beneficial in increasing the intelligibility of poor dialogue recordings.

The method used to suppress reverberation is quite different from that applied in the previous cases.

First, identify the range in which the reverberation lies

In general, reverberant spaces include soft materials that absorb high frequencies more rapidly than middle and lower frequencies. Even bare rooms with hard walls include these materials; they are the people who are speaking. Consequently, you will find that Low+Mid is often the most appropriate combination of ranges for suppressing reverberation.

Second, set the Band Gain controls

Set all six Band Gain controls to -24dB. This will ensure that (if the Level is set correctly) the DNS One applies maximum suppression to the tails of the sound.

Third, optimise the Level control

Starting at its minimum position, increase the Level control slowly. At some point before full noise suppression becomes apparent, you will hear the reverb tails of louder sounds become truncated. You can adjust the amount of truncation using the Level control rather than the Band Gain controls.

The final configuration might look like this:

Controls

Range Selectors

Selecting one of these concentrates all of the filters within the filter bank across that part of the spectrum.

Low	20Hz - 400Hz
Mid	200Hz - 6kHz
High	4kHz - 18kHz
Low+Mid	20Hz - 6kHz
Mid+High	200Hz - 18kHz
Full range	20Hz - 18kHz

Level

This tells the DNS One how much noise is present in the input. This figure is displayed in decibels in the numeric field.

Band Gain Controls

The six faders represent six frequency bands distributed from lower frequencies (left) to higher frequencies (right) across the selected range, and determine the maximum amount of processing that the DNS One will apply in each band. The maximum attenuation in each is displayed in decibels in its numeric field.

Adjust these to suppress as much noise as possible without introducing unwanted artefacts into the desired audio.

The bar graphs offer a visual indication of the activity in each of the Bands.

The  icons move all six Band Gain controls to 0dB (upper button) or to -24dB (lower button).

Learn

When On, DNS One calculates an estimate of the Level and determines suitable Band Gains.

ScreenVox

Near-zero latency noise reduction for speech

ScreenVox is a further development of CEDAR's Academy Award winning dialogue noise suppression (DNS) technology. Using it can be as simple as adjusting the attenuation, or you can use the Focus and Bias controls to fine-tune the process and obtain the desired result. Optimised for speech, it uses CEDAR's near-zero latency Quantum technology, which means that it can be used in all live sound and live-to-air scenarios.

Tutorial

In many cases, the default values for the Focus and Bias will be suitable, so using ScreenVox can be as simple as adjusting the Attenuation to obtain the desired amount of noise reduction.

If the noise level is particularly high, the output may benefit from increasing the Bias a little. This will identify more noise, but you should be careful that you do not make the signal sound too dry. If you wish to retain more air in the output, you should decrease the Bias, but be careful that you do not reintroduce unwanted noise into the signal.

If the noise lies in a narrow band, you can increase the Focus to allow the algorithm to concentrate on the dominant frequencies. If the noise is evenly distributed throughout the spectrum, you can reduce the Focus so that all frequencies are treated equally.

If necessary, experiment with the Bias and Focus to achieve the best result, and then fine-tune the Attenuation.

If you feel that ScreenVox is not removing the noise correctly or if the nature of the noise has changed suddenly, you can press Reset to reinitialise its internal parameters. It will then relearn the noise in the signal from that point onward.

Controls

	**Focus** This determines how ScreenVox responds to noise elements of greater or lesser prominence. **Focus = 100** Concentrates the processing upon the most prominent elements of the noise. **Focus = 0** All noise elements are treated equally. If the signal contains whines and buzzes, you can use a higher Focus to suppress these.
	**Bias** The Bias control allows you to bias the algorithm toward greater or lesser identification of the noise. **+ve values:** If more of the signal is determined to be noise, some of the fine detail in the wanted signal may be lost, but more noise will be removed. **-ve values:** If less of the signal is determined to be noise, more of the fine details of the wanted signal will be retained, but at the expense of reduced noise reduction.
	**Attenuation** Once the other parameters have been determined, whether at their default settings or following adjustment, the Attenuation controls the amount of noise removed.
	**Reset** Clears the internal parameters within the process, allowing it to start learning the noise afresh.

StageVox

Near-zero latency noise reduction for singers

StageVox incorporates unique algorithmic features that make it ideally suited to removing noise from singing voices. Using it can be as simple as adjusting the attenuation, or you can use the Focus and Ambience controls to fine-tune the process and obtain the desired result. Optimised for singing, it uses CEDAR's near-zero latency Quantum technology, which means that it can be used in all live sound and live-to-air scenarios.

Tutorial

In many cases, the default values for the Focus and Ambience will be suitable, so using StageVox can be as simple as adjusting the Attenuation to obtain the desired amount of noise reduction.

If the ambient noise is particularly high, the output may benefit from reducing the Ambience a little, but be careful that you do not make the signal sound too dry. If you wish to retain more air, you should increase the Ambience, but be careful that you do not reintroduce unwanted noise.

If the noise lies in a narrow band, you can increase the Focus to allow the algorithm to concentrate on the dominant frequencies. If the noise is evenly distributed throughout the spectrum, you can reduce the Focus so that all frequencies are treated equally.

If necessary, experiment with the Ambience and Focus to achieve the best result, and then fine-tune the Attenuation.

If you feel that StageVox is not removing the noise correctly or if the nature of the noise has changed suddenly, you can press Reset to reinitialise its internal parameters. It will then relearn the noise in the signal from that point onward.

Controls

	**Focus** This determines how StageVox responds to noise elements of greater or lesser prominence. **Focus = 100** Concentrates the processing upon the most prominent elements of the noise. **Focus = 0** All noise elements are treated equally. If the signal contains whines and buzzes, you can use a higher Focus to suppress these.
	**Ambience** Lower the Ambience to decrease the amount of ambience in the output and make the signal sound drier. Increase the Ambience to increase the amount of ambience in the output. This will 'open' the sound, but may remove less noise.
	**Attenuation** Once the other parameters have been determined, whether at their default settings or following adjustment, the Attenuation controls the amount of noise removed.
	**Reset** Clears the internal parameters within the process, allowing it to start learning the noise afresh.

VoicEX 2

VoicEX 2 is based upon CEDAR's latest advances in Artificial Intelligence and Machine Learning. It incorporates some unique algorithmic features that take it beyond standard DNN (deep neural network) based noise reduction and, for a wide range of material, is able to separate voices (primarily but not limited to speech) from other sounds, and to perform superior noise reduction when compared with traditional signal processing techniques.

Tutorial

Internally, VoicEX 2 creates two signal streams - one containing any voices identified within the input, and the other containing everything else. To eliminate noise and excessive reverb, leave the Voice control at zero and reduce the Noise to the desired degree. Do not over process; although VoicEX 2 is extremely tolerant, you may introduce a change in the retained signal. To suppress the voices, leave the Noise knob at zero and reduce the Voice. You can accentuate the voices or the background by increasing either knob above zero.

Controls

	**Voice** Attenuate or accentuated signal that is identified as voices. The range is notionally -50dB to +20dB.
	**Noise** Attenuate or accentuated signal that is not identified as voices. The range is notionally -50dB to +20dB.

Licence and Limited Warranty

1. DEFINITIONS

In this Licence and Limited Warranty the following words and phrases shall bear the following meanings:

'the Company' is CEDAR Audio Limited of 20 Home End, Fulbourn, Cambridge, CB21 5BS, UK;

'the System' means any instance of the CEDAR products developed by the Company, including but not limited to Adaptive Limiter 2, Declick, Declip, Decrackle, DNS One, ScreenVox, StageVox, and VoicEX 2;

'this Document' means this Licence and Limited Warranty.

2. ISSUE AND USE OF THE SYSTEM

2.1 The terms and conditions of this Document are implicitly accepted by any person or body corporate who shall at any time use or have access to the System, and are effective from the date of supply of the System by CEDAR Audio Limited to its immediate customer.

2.2 The Company hereby grants to the Licensee and the Licensee agrees to accept a non-exclusive right to use the System.

3. PROPERTY AND CONFIDENTIALITY

3.1 The System contains confidential information of the Company and all copyright, trademarks, trade names, styles and logos and other intellectual property rights in the System including all documentation and manuals relating thereto are the exclusive property of the Company. The Licensee acknowledges that all such rights are the property of the Company and shall not question or dispute the ownership of any such rights nor use or adopt any trading name or style similar to that of the Company.

3.2 The Licensee shall not attempt to reverse engineer, modify, copy, merge or transcribe the whole or any part of the System or any information or documentation relating thereto.

3.3 The Licensee shall take all reasonable steps to protect the confidential information and intellectual property rights of the Company.

4. LIMITED WARRANTY AND POST-WARRANTY OBLIGATIONS

4.1 The Company warrants that the System will perform substantially in accordance with the appropriate section of its accompanying product manual for a period of one year from the date of supply to the Company's immediate customers.

4.2 The Company will make good at its own expenses by repair or replacement any defect or failure that develops in the System within one year of supply to the Company's immediate customer.

4.3 The Company shall have no liability to remedy any defect, failure, error or malfunction that arises as a result of any improper use, operation or neglect of the System, or any attempt to repair or modify the System by any person other than the Company or a person appointed with the Company's prior written consent.

4.4 In the case of any defect or failure in the System occurring more than twelve months after its supply to the Company's immediate customer the Company will at its option and for a reasonable fee make good such defect or failure by repair or replacement (at the option of the Company) subject to the faulty equipment having first been returned to the Company. The Company will use reasonable efforts to return repaired or replacement items promptly, all shipping, handling and insurance costs being for the account of the Licensee.

4.5 The above undertakings 4.1 to 4.4 are accepted by the Licensee in lieu of any other legal remedy in respect of any defect or failure occurring during the said period and of any other obligations or warranties expressed or implied including but not limited to the implied warranties of saleability and fitness for a specific purpose.

4.6 The Licensee hereby acknowledges and accepts that nothing in this Document shall impose upon the Company any obligation to repair or replace any item after a time when it is no longer produced or offered for supply by the Company or which the Company certifies has been superseded by a later version or has become obsolete.

5. FORCE MAJEURE

The Company shall not be liable for any breach of its obligations hereunder resulting from causes beyond its reasonable control including, but not limited to, fires, strikes (of its own or other employees), insurrection or riots, embargoes, container shortages, wrecks or delays in transportation, inability to obtain supplies and raw materials, or requirements or regulations of any civil or military authority.

6. WAIVER

The waiver by either party of a breach of the provisions hereof by the other shall not be construed as a waiver of any succeeding breach of the same or other provisions, nor shall any delay or omission on the part of either party to exercise any right that it may have under this Licence operate as a waiver of any breach or default by the other party.

7. NOTICES

Any notices or instruction to be given hereunder shall be delivered or sent by first-class post or telecopier to the other party, and shall be deemed to have been served (if delivered) at the time of delivery or (if sent by post) upon the expiration of seven days after posting or (if sent by telecopier) upon the expiration of twelve hours after transmission.

8. ASSIGNMENT AND SUB-LICENSING

The Licensee may at his discretion assign the System and in doing so shall assign this Licence its rights and obligations to the purchaser who shall without reservation agree to be bound by this Licence. The original Licensee and any subsequent Licensees shall be bound by the obligations of this Licence in perpetuity.

9. LIMITATION OF LIABILITY

The Company’s maximum liability under any claim including any claim in respect of infringement of the intellectual property rights of any third party shall be, at the option of the Company either:

(a) return of a sum calculated as the price received for the System by the Company from its immediate customer depreciated on a straight line basis over a one year write-off period; or

(b) repair or replacement of those components of the System that do not meet the warranties contained within this Document.

The foregoing states the entire liability of the Company to the Licensee.

10. CONSEQUENTIAL LOSS

Even if the Company has been advised of the possibility of such damages, and notwithstanding anything else contained herein the Company shall under no event be liable to the Licensee or to any other persons for loss of profits or contracts or damage (whether direct or consequential) arising in connection with the System or any modification, variation or enhancement thereof and including any documentation or data provided by the Company or for any other indirect or consequential loss.

11. ENTIRE AGREEMENT

The Company shall not be liable to the Licensee for any loss arising in connection with any representations, agreements, statements or undertakings made prior to the date of supply of the System to the Licensee.

12. TERMINATION

This Licence may be terminated forthwith by the Company if the Licensee commits any material breach of any terms of this Licence. Forthwith upon such termination the Company shall have immediate right of access to the System for the purpose of removing it.

13. SEVERABILITY

Notwithstanding that the whole or any part of any provision of this Document may prove to be illegal or unenforceable the other provisions of this Document and the remainder of the provision in question shall remain in full force and effect.

14. HEADINGS

The headings to the Clauses are for ease of reference only and shall not affect the interpretation or construction of this Document.

15. LAW

This Document shall be governed by and construed in accordance with English law and all disputes between the parties shall be determined in England in accordance with the Arbitration Act 1950 and 1979.

CEDAR Audio Limited
20 Home End
Fulbourn
Cambridge
CB21 5BS
United Kingdom

t: +44 1223 881771
e: info@cedaraudio.com
w: www.cedaraudio.com