

Dialog 2

User Manual



Last updated: July 16, 2020
Copyright © 2020, Wave Arts, Inc. All Rights Reserved

Table of Contents

1. Dialog Overview.....	4
1.1 Audio Routing	4
1.2 Using Dialog	6
2. Brickwall Filters.....	8
2.1 About Brickwall Filters	8
2.2 Using the Brickwall Filters	8
3. De-Hum.....	10
3.1 About Hum	10
3.2 About Buzz	10
3.3 Using De-Hum.....	10
4. De-Noise.....	12
4.1 About Noise	12
4.2 Noise Reduction.....	12
4.4 Using De-Noise.....	13
4.5 Learning the Noise Floor	13
4.6 De-noise Spectrum Display.....	14
4.7 Reduction Amount	15
4.8 De-noise Presets.....	15
5. De-Plode.....	16
5.1 About Plosives.....	16
5.2 Using De-plode.....	16
6. De-Ess	19
6.1 Sibilants and De-essing.....	19
6.2 Using De-Ess.....	19
7. Equalizer	22
7.1 Equalization	22
7.4 Notch filter	24
7.5 Bandpass filter	25
7.6 Lowpass and Highpass Filters.....	26
7.7 Using the Equalizer	27
8. Compressor.....	30
8.1 About Compressors	30
8.2 Using the Compressor	31
9. Output and Limiter	33
10. Installation and Registration.....	34
10.1 Installation.....	34
10.2 Registration	34
10.3 Wave Arts licensing	34
10.4 PACE/iLok licensing	34
11. Plug-in Control Operation.....	36
11.1 Knobs	36
11.2 Text Entry.....	36
11.3 Selector button.....	37
11.4 Sliders.....	38
11.5 Buttons	38

11.6 Output Meters	38
12. Menu Bar and Preset Manager	39
12.1 Bypass	39
12.2 Undo.....	39
12.3 Copy	39
12.4 A/B buffers	39
12.5 Preset name and arrow controls	40
12.6 Preset menu.....	40
12.7 Factory Presets.....	40
12.8 User Presets.....	41
12.9 Save As.....	41
12.10 Save	41
12.11 Import... ..	41
12.12 Export... ..	42
12.13 User Preset Manager.....	42
12.14 Tools menu	42
12.15 About.....	43
12.16 Register.....	43
12.17 Offline Register.....	44
12.18 Open User Manual... ..	44
12.19 Preferences.....	44
License Agreement.....	45
Support.....	47
Trademark Notices	48
Index.....	49

1. Dialog Overview



Figure 1-1. Dialog user interface.

The Dialog plug-in is intended for post production clean-up and processing of the spoken voice. Because it contains many general purpose functions it can be used for a wide range of clean up and processing applications. Dialog's key features are listed below:

- Low and high pass brickwall filters with various rolloff slopes
- Hum and buzz removal
- Broadband noise removal
- De-plodge function for plosive or wind rumble attenuation
- De-ess function for sibilant attenuation
- 10-band EQ section with 11 different filter types and 64-bit processing
- Compressor
- Lookahead peak limiter

1.1 Audio Routing

Dialog's audio routing and meter placement is shown in the diagram below:

3. Menu Bar and Preset Manager

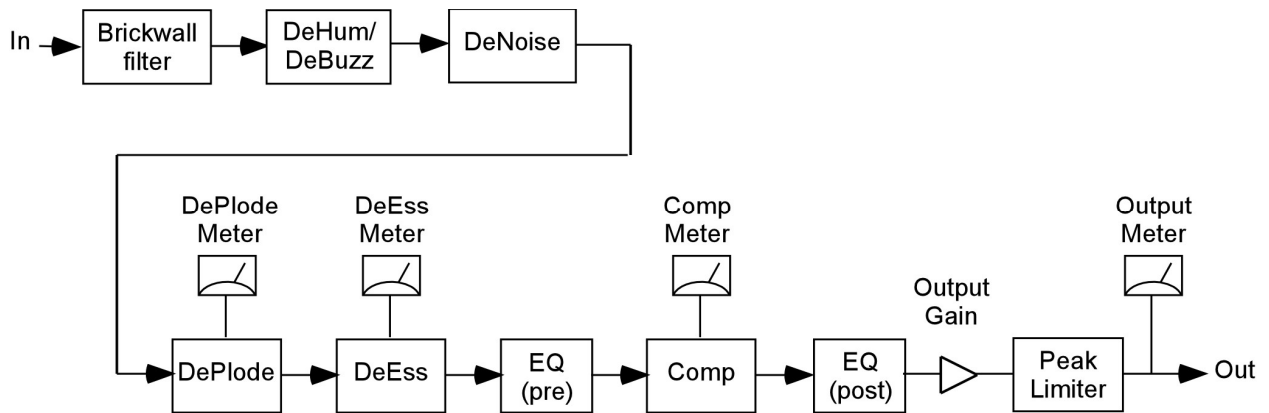


Figure 1-2. Dialog audio routing.

The input signal is processed first by the brickwall filter. This allows the user to eliminate noise signals that are outside the frequency range of the recorded speech. Thumps and rumble can be eliminated using a highpass brickwall, while high frequency hiss or tones can be eliminated using a lowpass brickwall.

The signal is next processed by the hum and buzz remover. The hum remover notches 10 harmonics of either 50 Hz or 60 Hz hum. For really bad hums that sound like a buzz, the buzz remover can be selected, again for either 50 Hz or 60 Hz line frequency. The buzz remover imparts a hollow coloration (like a comb filter) to the sound and should only be used as a last resort (try using just the denoiser first).

The signal is processed next by the noise remover which employs a sophisticated algorithm for attenuating background noises. The noise remover acts like a bank of hundreds of noise gates each operating in a different frequency range. In order to set the threshold level for all the noise gates, the noise profile must be learned by processing a short portion of background noise as described later in this manual. After the noise profile is learned it is important not to change the settings for the brickwall filter or the hum and buzz remover. Doing so will change the noise profile and will require that you re-learn the noise profile.

The deploder runs after the denoiser. The deploder is used to attenuate loud plosives which can occur when a talker is too close to the microphone. It can also attenuate rumble due to wind hitting the microphone. The Dialog deploder is a side chain compressor with a lowpass filter in the sidechain path. Signals that contain bursts of low frequency energy cause the compressor to engage and attenuate the entire signal during the burst.

The deesser runs after the deploder. The deesser is used to attenuate loud sibilant ("ess") sounds which depend on the talker's voice and the microphone setup. The Dialog deesser works by splitting the signal into low

and high frequency components and then applying a compressor to the high frequency component. High frequency energy due to a sibilant causes the compressor to engage and attenuate the high frequencies.

Equalization (EQ) is the principal tool to shape the tonal character of the speech. If the EQ routing is set to "pre", the EQ runs after the deesser and before the compressors; if set to "post", it runs after the compressors.

The compressor runs next. The compressor is used to even out variations in loudness which can occur when a single talker is too close to the microphone, or when multiple talkers are at different distances from the microphone.

If the EQ routing is set to "post", the EQ runs after the compressor.

The output gain and peak limiter run last. The peak limiter can be used as a hard compressor / loudness maximizer, by cranking up the output gain, or used simply to prevent peaks from exceeding -0.1 dB.

All of the above functions are described in more detail in later sections.

1.2 Using Dialog

Dialog contains eight different audio processors useful for cleaning up, equalizing, and compressing speech audio. By combining all the processors into a single plug-in, only one instance of Dialog is needed. However, one can also use multiple instances of Dialog, with different sections enabled. Remember that only the sections that are enabled consume CPU capacity.

The audio routing described in the previous section also suggests the proper order to apply processing:

- 1) Start by cleaning up hum and noise. In particular, apply brickwall filters first to eliminate out of band signals, then apply hum/buzz filters to eliminate hum and buzz, and then learn the noise floor and apply noise reduction. Remember that once you learn the noise floor and apply noise reduction, any changes to the hum/buzz filters or brickwall filters will require that you relearn the noise floor.
- 2) After the audio is cleaned up, apply de-plode if needed to attenuate plosives or wind rumble.
- 3) Apply de-ess to tame loud sibilants, if needed.
- 4) Apply compression and equalization to tailor the sound to suit. You may want to compress first and then equalize or the other way around. Also the equalizer can run before or after the compressor in the signal chain depending on the pre/post setting.

3. Menu Bar and Preset Manager

5) Adjust output levels and peak limit if maximizing loudness.

The following sections describe the various processors within Dialog. Each section starts with a general description of the audio topic, and then a specific discussion of how to use the corresponding Dialog processor.

2. Brickwall Filters

2.1 About Brickwall Filters

Brickwall filters are lowpass or highpass filters with very steep cutoffs, used to pass all frequencies up to the cutoff frequency and eliminate all frequencies beyond the cutoff. The frequency range that is passed unaltered is called the "passband", the frequency range that is attenuated is called the "stopband".

Dialog has several different types of brickwall filters with different rolloff slopes, listed in the table below:

Rolloff slope	Filter Type
60 dB / octave	10 th order elliptical
48 dB / octave	8 th order Butterworth
36 dB / octave	6 th order Butterworth
24 dB / octave	4 th order Butterworth

Elliptical filters are capable of very steep rolloffs, but have a slight amount of ripple in the passband and stopband. Dialog's 60 dB/octave brickwall filters are implemented using 10th order elliptical filters, which have at least 90 dB of stopband attenuation and less than 0.1 dB of passband ripple. The other filters with shallower rolloffs are implemented using Butterworth filters which have no ripple.

Brickwall filters are used to eliminate unwanted frequency ranges. Typically, a brickwall filter would be used when processing a noisy recording of an instrument sound that does not use the entire frequency range. The brickwall filter would be positioned at the edge of the instrument's frequency range to eliminate out-of-band noise. So for example, when processing voice, one could use the brickwall filters to eliminate all frequencies below 100 Hz and above 8 kHz.

The brickwall filters can also be used to zero in on a particular frequency range just for analysis purposes.

2.2 Using the Brickwall Filters



Figure 2-1. Brickwall filter user interface.

3. Menu Bar and Preset Manager

The user-interface of the brickwall filters is shown above. At the left is a selector button to choose the rolloff slope of the filters. At the right are enable buttons to enable/disable the hignpass and lowpass filters, and knobs to adjust the cutoff frequencies. When either band is enabled, a corresponding brickwall control handle appears in the EQ frequency response display as shown below. You can change the brickwall frequency by dragging the control handle.




Figure 2-2. Brickwall control handles in the EQ display.

3. De-Hum

3.1 About Hum

Hum is caused by interference of audio circuitry with line power. Line power oscillates at either 50 Hz or 60 Hz depending on your world location. So in its simplest form, hum is a low frequency interference tone at either 50 Hz or 60 Hz. However, Hum often contains harmonics (multiples) of the fundamental frequency. For example, in a region with 60 Hz power, interfering hum would typically contain frequencies of 60 Hz, 120 Hz, 180 Hz, and so on.

Hum is best removed by using a set of sharp notch filters to eliminate the fundamental and harmonic frequencies of the hum. Sharp notch filters work well because the frequency of line power is very stable so the filters can be very narrowly tuned. Dialog uses 10 notch filters to eliminate up to the 10th harmonic.

The frequency of hum depends on the frequency of the local line power. 60 Hz power is used in North America (except Greenland), Central America, most of South America, Saudi Arabia, and parts of southeast Asia including Korea. The rest of the world, including Europe, Asia, Africa, and Australia uses 50 Hz power.

3.2 About Buzz

Buzz is a periodic impulsive contamination which like hum is usually caused by poor grounding. We define the term buzz to mean a hum that contains a lot of high frequency energy. Consequently, buzz contains harmonics of the fundamental that extend through the entire audible frequency range.

Buzz is eliminated by applying a set of notch filters located at the fundamental frequency and at all harmonics in the audio range. The set of notch filters is created using a "comb filter" which essentially computes the difference of the input signal with a copy of the input signal delayed by one fundamental period. The difference calculation removes any periodic contamination. However, the side effect imparted to the signal is that of a comb filter, that is, an echo with a very short delay. This can sound like a hollow room ambience and may be objectionable. Hence, buzz removal imparts significant side effects.

3.3 Using De-Hum



3. Menu Bar and Preset Manager

Figure 3-1. De-hum/de-buzz user interface.

The Dialog De-Hum interface, shown above, is extremely simple. The enable button enables/disables the processing. The mode, either hum or buzz removal, is selected by clicking on the appropriate button. When "hum" is selected, the first 10 harmonics are notched; then "buzz" is selected, all audible harmonics are notched, which can create a hollow echo sound. The fundamental frequency, either 50 or 60 Hz, is selected by clicking on the appropriate button.

Because Dialog is designed for post production, we assume the audio is being processed from an original recording, likely a digital recording, and hence the fundamental frequency will not have changed due to recording and playback. When processing recordings from magnetic tape or vinyl, the playback may not be at the same speed as the recording and hence the fundamental frequency will have shifted. In this case you may need to use a different processor, such as MR Hum, which can accept a continuous range of fundamental frequencies.

4. De-Noise

4.1 About Noise

Noise is a general term that can apply to various types of unwanted artifacts in an audio signal, including hum, buzz, clicks, crackle, rumble, and hiss. We use the term “broadband noise” to describe noise consisting of a random signal which sounds like hiss. This type of noise is present in every audio signal. There are many sources of broadband noise:

Digital quantization – in digital signals, the quantization of the signal to a certain bit resolution creates broadband noise called quantization noise.

Acoustic recordings – there is always some acoustic noise, which can be due to ventilation sounds or other background noises.

Electronics – all electronics, particularly high gain amplifiers, will add noise to a signal due to thermal variation in the amplifier elements.

Magnetic tape recordings – magnetic tape is composed of individual grains of magnetic material which are each magnetized differently to record an audio signal. The finite number of grains and limited tape speed causes random noise during playback.

Phonograph recordings – surface irregularities cause playback noise.

An important characteristic of the noise is its spectrum, that is, the distribution of energy at different frequencies. The noise spectrum is also called the noise “floor”, because it defines the lowest signal level at each frequency that can be distinguished from the noise. In order for a signal to be audible, it has to be above the noise floor. The noise floor is also called the noise “profile” in other noise reduction literature.

For many sources of noise, the noise floor spectrum will remain fairly constant over time. This is typical of quantization noise, electronic noise, magnetic tape noise, and phonograph noise. Acoustic noise, however, can either be constant or varying. We would expect the acoustic noise in a quiet recording room to be fairly constant. But the acoustic noise in an outdoor environment can vary significantly over time due to traffic noise, wind, or other background sound.

4.2 Noise Reduction

The noise reduction algorithm used by Dialog is very effective at removing noise with a constant spectrum. The Dialog De-noise algorithm operates in the frequency domain. At each frequency, it compares the signal level with the noise floor. If the signal level is well above the noise floor, it is unchanged, but signals that are close to the noise floor are attenuated. This is like having a noise gate running at each frequency, where the gate threshold is set to the noise floor level at that frequency.

3. Menu Bar and Preset Manager

In order to determine the noise floor spectrum, the input signal spectrum is sampled and successive samples are averaged together. This process is typically called "learning". The user is responsible for selecting a portion of noise to analyze and this noise is processed while the algorithm is in learning mode. The Dialog algorithm requires only a short noise sample, down to 50 msec, to learn the noise spectrum effectively, although longer learning times give more accurate results.

To handle varying noise floors, there are two approaches. One can learn a long varying portion of noise and use this average floor, or one can segment the audio into different clips and learn different noise floors for each clip. The latter method is more effective.

4.4 Using De-Noise



Figure 4-1. De-noise user interface.

The De-noise user interface is shown above. In addition to the above controls there is a graphical display shared with the EQ display. In order to select the De-Noise display, click on the DE-NOISE tab under the display. For stereo signals, the DISP MODE control selects whether the display shows the left channel, right channel or both channels. Clicking on the vertical axis shows a popup menu of axis scaling choices.

4.5 Learning the Noise Floor

The first step is to identify a section of your track that contains only background noise. Often this can be found at the very start of the track before the signal comes in. However, if the track has already been cropped to eliminate leading and trailing signal, then you will have to find noise during a silent section in the track. Dialog can learn a noise profile with as little as 50 msec of noise, so this could be found between words in a dialog track for example.

The learn time is specified using the LEARN TIME parameter. It has various choices from 50 msec to 5 seconds, and also has a manual mode. Choose the longest learn time that is shorter than the portion of noise you have identified. In manual mode, learning begins when you click the LEARN button and ends when the LEARN button is clicked again. All the other modes learn for the specified amount of time after the LEARN button is pressed.

After choosing the learn time, position your editor's audio transport so that audio playback will commence at the noise segment, then click LEARN to enable learning. The LEARN button will glow with a purple color. Then start playback. The LEARN button will extinguish when learning is complete. Note that clicking LEARN will automatically enable the De-noise processor if it is not already enabled.

4.6 De-noise Spectrum Display

During learning the display shows the dynamic spectrum of the signal in red, with the accumulated average noise floor drawn in yellow, as shown in the figure below. When learning is complete the yellow noise floor stops changing and noise reduction processing begins.



Figure 4-2. De-noise frequency display during learning. Mono display mode is selected.

While noise reduction is in progress, the input signal spectrum is drawn in green and the output spectrum is drawn in red. The figure below shows an example display of a single channel during noise reduction. Noise reduction is occurring at any frequency where the red contour is below the green contour, shown with an orange fill.



Figure 4-3. De-noise frequency display during processing. Mono display mode is selected.

4.7 Reduction Amount

It is often necessary to limit the maximum amount of noise attenuation. This is done via the Amount parameter, which specifies the maximum attenuation Dialog will apply. By limiting the maximum attenuation, we limit the amount of noise reduction, but also limit processing artifacts.

4.8 De-noise Presets

The preset control allows you to save the learned noise floor as a user preset. This can be useful if you need to save the noise floor to apply to a different track. To save the learned noise floor as a user preset, click on the De-noise preset button and select "Save As..." from the popup menu. This will ask you for a preset name. The saved preset will appear in the De-noise preset popup menu in the "User" submenu.

Some audio editors do not let you change the audio selection while a plug-in is being used. In this case, you will need to select the noise segment first, then open Dialog and learn the noise floor, then save the floor as a user preset, close Dialog, select the entire track (or cancel the existing selection), open Dialog again, select the saved noise floor preset, and then process the entire track.

5. De-Plode

5.1 About Plosives

Plosives are consonant sounds created by the sudden opening of the oral passage accompanied by a burst of air. The consonants [p], [t], [d] are examples. When a speaker's mouth is close to a microphone, the bursts of air may cause large excursions of the microphone element, recorded as high amplitude low frequency noise. A similar effect happens when wind hits a microphone. Without the use of a wind shield the wind causes low frequency rumble or booms to be recorded.

To remove the low frequency impulses caused by close miking or wind we can use a side chain compressor where the side chain equalizer passes low frequencies, such as below 200 Hz. Hence, low frequency bursts cause the compressor to engage and attenuate the signal. This works better than a high pass filter because the plosive bursts also contain high frequency energy which would not be attenuated by a high-pass filter. Also, the signal is not affected unless there is a burst of low frequency energy.

The Dialog De-plode processor is based on a side-chain compressor with a lowpass filter (LPF) in the side-chain. The audio routing is shown in the figure below.




Figure 5-1. De-plode audio routing. When the monitor button is enabled the side-chain signal is output.

5.2 Using De-plode

The de-plode user interface is shown in the figure below.

3. Menu Bar and Preset Manager



Figure 5-2. De-plode user interface.

Clicking ENABLE enables the de-plode processor. For most deploding needs, start with one of the factory presets such as Basic Deplode. Clicking on the preset selector button will cycle through the factory (and any user) presets. Clicking on the preset name will show a popup menu of preset options.

The meters are below the preset control. On the left is an input meter showing the level of the side-chain signal, i.e. after processing by the lowpass filter. The peak input level to date is shown at the top of the meter; clicking the value will reset it. On the right is the attenuation meter showing the attenuation of the compressor. The peak attenuation is shown at the top of the meter; clicking the value will reset it. The range of the attenuation meter can be changed by right-clicking (SHIFT-clicking on Mac) the meter axis values, which shows a popup menu of meter ranges.

There are standard compressor controls for threshold, reduction, attack and release times. The threshold determines the input level that must be exceeded for the compressor to attenuate. The threshold can be modified using the knob control or the slider next the input meter. The reduction control is a standard compressor ratio, it determines how much attenuation is applied as the input exceeds the threshold. The attack time determines how fast the compressor will attenuate when the side-chain input exceeds the threshold value. For most deploding you will want to use fast times such as 1 msec. The release time determines how quickly the compressor restores the gain after the gain has ducked. For more information about compressor controls, see the section on the Dialog Compressor.

The side-chain lowpass filter cutoff is controlled by the SC FREQ control. Selecting lower frequencies will reduce the input level and selecting higher

frequencies will increase the input level. To hear the side-chain signal, i.e. the output of the lowpass filter, click the MONITOR button.

6. De-Ess

6.1 Sibilants and De-essing

Esses, or sibilants, refer to the consonant sounds created by forcing air through a constricted passage. The resulting turbulence of air creates a high frequency hissing sound. The ess sounds “ss” and “sh” are typically the loudest sibilants. With close-microphoned dialog recordings, it is easy for the sibilants to be recorded too loud. A de-esser processor reduces the level of the esses to make the dialog more natural sounding.

There are various ways to build a de-esser, usually using a side-chain compressor with a highpass or bandpass filter in the side-chain. The filter is tuned to respond to sibilant frequencies, generally from 4 kHz to 10 kHz, such that when sibilants are detected the compressor ducks the gain. Using a compressor is better than simply applying a lowpass or shelving equalizer to reduce high frequencies, because an equalizer will also reduce high frequencies in the non sibilant speech, causing the speech to sound dull or muffled.

The Dialog de-esser uses a multiband approach. The signal is first split into low and high frequency bands using a linear-phase bandsplitter. The use of a linear-phase bandsplitter allows the input signal to be perfectly reconstructed by summing the low and high bands, without phase or amplitude distortion. The high frequency band is processed by a compressor which ducks the gain when the high frequency signal exceeds a threshold. The audio routing of the Dialog de-esser is shown in the figure below. This architecture is like a side-chain compressor with a high pass filter in the side-chain, except the gain ducking acts like a shelving equalizer, i.e., only the high frequencies are ducked rather than all frequencies.

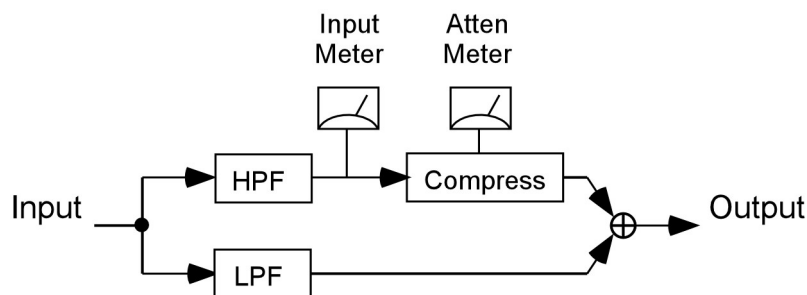


Figure 6-1. De-ess audio routing. When the monitor button is enabled the high-pass filter (HPF) is output.

6.2 Using De-Ess

The Dialog de-ess user interface is shown in the figure below.



Figure 6-2. De-ess user interface.

Clicking ENABLE enables the de-ess processor. For most de-essing needs, start with one of the factory presets such as Mild Lo. Clicking on the preset selector button will cycle through the factory (and any user) presets. Clicking on the preset name will show a popup menu of preset options. The “Mild” presets apply a mild amount of reduction while the “Deep” presets have more reduction. The “Lo” presets have a lower sibilant frequency while the “Hi” presets have a higher sibilant frequency.

The meters are below the preset control. On the left is an input meter showing the level of the high frequency band signal, i.e. after processing by the highpass filter. The peak input level to date is shown at the top of the meter; clicking the value will reset it. On the right is the attenuation meter showing the attenuation of the compressor. The peak attenuation is shown at the top of the meter; clicking the value will reset it. The range of the attenuation meter can be changed by right-clicking (SHIFT-clicking on Mac) the meter axis values, which shows a popup menu of meter ranges.

There are standard compressor controls for threshold, reduction, attack and release times. The threshold determines the input level that must be exceeded for the compressor to attenuate. The threshold can be modified using the knob control or the slider next the input meter. The reduction control is a standard compressor ratio, it determines how much attenuation is applied as the input exceeds the threshold. The attack time determines how fast the compressor will attenuate when the side-chain input exceeds the threshold value. The release time determines how quickly the compressor restores the gain after the gain has ducked. For more information about compressor controls, see the section on the Dialog Compressor.

3. Menu Bar and Preset Manager

The lowpass/highpass crossover frequency is controlled by the SC FREQ control. Because sibilant energy is generally concentrated at frequencies from 4 kHz to 6 kHz, a crossover frequency of 4 kHz usually suffices. Selecting higher frequencies will slightly reduce the input level and selecting lower frequencies will slightly increase the input level. To hear the side-chain signal, i.e. the output of the highpass filter, click the MONITOR button.

7. Equalizer

7.1 Equalization

The equalizer is used to boost certain frequencies and reduce others, thus changing the tonality of a sound. It is helpful to have a general understanding of the frequency ranges in a mix:

Low Bass (20 – 60 Hz)- controls rumble, this range is best reproduced on a subwoofer equipped system.

Bass (60 to 250 Hz)- the low end of the mix, where the fundamentals of bass and other rhythm oriented instruments like the kick drum reside. Controls the overall fullness/roundness of a track or mix. Reducing around the 100 to 200 Hz range can help reduce "boominess".

Low Midrange (250 to 2000 Hz)- many of the low harmonics of most instruments are in this range, and some boost here between 250 and 800 Hz can improve clarity of lower pitched instruments. Too much boost in this whole range can lend to a telephone-like quality (boost from 500Hz to 1Khz can sound horn-like, from 1Kz to 2Kz can sound tinny), and often frequencies in this area can be reduced on mid-range instruments like guitar, vocals and keyboards to improve a mix.

High Midrange (2000 Hz – 4000 Hz)- this range is an important speech recognition area, and also determines projection and clarity of mid-range instruments. Too much boost in this area can be fatiguing on the ear.

Presence (4000 Hz – 6000 Hz)- controls how close and distinct instruments and vocals sound, too much in this range will cause harshness.

Brilliance (6000 Hz – 20 kHz)- this range is associated with clarity, "sizzle", and "air".

Dialog's equalizer provides up to 10 separate EQ bands, more than needed for most jobs. Each band has a choice of 11 different filter types per band. We've erred on the side of providing more control rather than skimping. All of the filter types are based on second order filters; this limits their rolloff slope to -12 dB per octave. The filter types are described below.

7.2 Parametric EQ

The parametric EQ type is one you will use often, as it allows you to reinforce or attenuate at a specific frequency point. The parametric EQ is defined in terms of the center frequency, the height of the boost/cut in dB, and the width. The width is measured in octaves between the half height points in the response. For example, if the center frequency is 2 kHz, the height is 12 dB,

and the width is 2 octaves, then the width is defined at the 6 dB points in the response, which will be at 1 kHz and 4 kHz (one octave below and one octave above the center frequency, respectively).




Figure 7-1. Parametric EQ with freq = 2000 Hz, height = 12 dB, and width = 2 octaves.

7.3 Shelf filter

The shelf filter boost or cuts by a fixed amount above or below a corner frequency. The shelf filter has parameters of corner frequency and shelf height in dB. The figure below shows a low shelf filter with a boost of 12 dB and a corner frequency of 200 Hz. The corner frequency is defined at the half height point in the response, so for example, in the figure below, the response is 6 dB at 200 Hz. Dialog provides three variations of the shelf filter which differ in the shape of the shelf transition: standard, resonant, and vintage. The standard shelf filter as shown in the figure has the steepest possible transition without having any overshoot. The "resonant" shelf filter has a variable transition slope defined by a resonance parameter. Filter resonances are usually defined using a Q parameter, called the "quality factor"; higher values of Q are more resonant and hence sharply tuned. Dialog uses the width knob to set the resonance, in units of Q. A shelf with a high resonance has a steep transition, but also overshoots symmetrically on each side, as shown in the figure below. The standard shelf filter is identical to the resonant shelf filter with a resonance (Q) of 0.707. Finally, the "vintage" shelf filter has a particular sort of asymmetrical overshoot which is quite gentle and pleasing sounding. The vintage shelf is modeled after the response shapes of certain analog equalizers.




Figure 7-2. Low shelf with freq = 1000, height = 12 dB.




Figure 7-3. Resonant low shelf with freq = 200, height = 12 dB, $Q = 2$.




Figure 7-4. Vintage low shelf with freq = 200, height = 12 dB, $Q = 1.414$.

7.4 Notch filter

The notch filter cuts a specific frequency. It has parameters of center frequency and width. The width of the notch is measured in octaves between the -3 dB points of the filter response. Dialog's notch filter eliminates the center frequency completely, this would correspond to a height of minus infinity dB. The notch is commonly used to eliminate hum caused by power line interference. In the US, power lines oscillate at 60 Hz, in Europe and other parts of the world power lines oscillate at 50 Hz. Overtones of the hum are common and require an additional notch filter to cancel each overtone. Another commonly seen interference tone is caused by the horizontal flyback oscillation of old analog televisions. In the US, NTSC TV sets oscillate at 15,734 Hz. In Europe, PAL and SECAM TV sets oscillate at 15,625 Hz. Often you will find these tones in older acoustic recordings made anywhere near TV sets.




Figure 7-5. Notch filter with freq = 60, width = 0.1 octaves.

7.5 Bandpass filter

The bandpass filter is the opposite of the notch, it passes only the frequencies in a band around the center frequency. The bandpass filter has parameters of center frequency and width. The width of the bandpass is measured in octaves between the -3 dB points of the filter response. The bandpass filter would typically only be used in sound design to simulate the sound of a reduced frequency response, such as a telephone. One could also use the brickwall filters for this. The bandpass filter rolls off at -12 dB per octave on each side of the center frequency, much gentler than the brickwall, which rolls off at 60 dB/octave. The bandpass filter can also be used in analysis to isolate a particular frequency. For example, if there is a contaminating tone in a signal, one can sweep a narrow bandpass filter to isolate the tone, then switch to a notch filter to eliminate the tone.




Figure 7-6. Bandpass filter with freq = 2 kHz, width = 1 octave.

7.6 Lowpass and Highpass Filters

Dialog also provides lowpass and highpass filters. These pass frequencies up to a corner frequency, then roll off at -12 dB per octave beyond. The lowpass and highpass filters have only one parameter: corner frequency. The corner frequency is defined at the -3 dB point in the response. The choice to use the lowpass and highpass filters or to use the brickwall filters depends on how steep you need to rolloff to be.




Figure 7-7. Lowpass filter with freq = 2 kHz.

7.7 Using the Equalizer

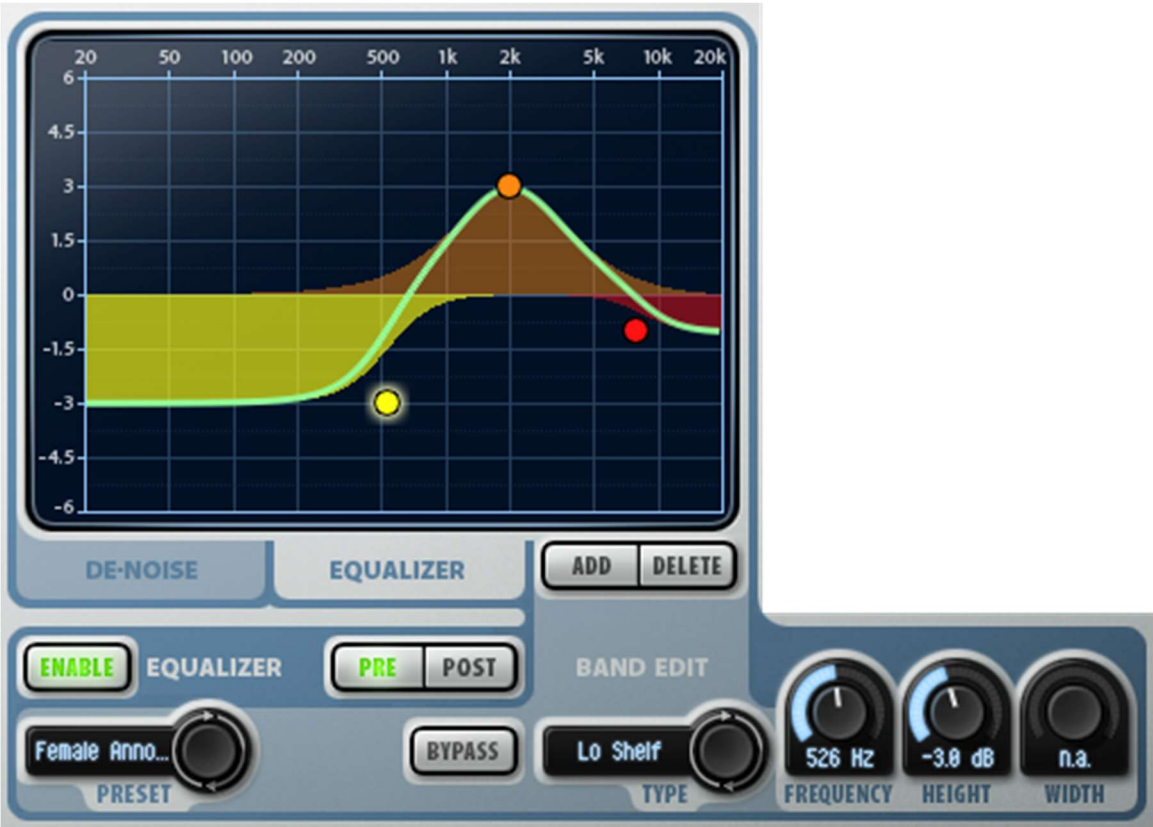


Figure 7-8. Dialog EQ User interface

The Dialog EQ interface is shown above. At the top is the frequency response display, which shows the frequency response of the current EQ. The frequency response display contains a number of control handles, drawn as colored balls, which correspond to EQ bands; clicking a handle will select that EQ band for editing, dragging a handle will change the parameters for the EQ. The band that is selected shows a halo around its control handle. The table below summarizes the operations that can be done inside the frequency response display:

EQ Section		
Control	Mac	Windows
Adjust an EQ band's frequency and height	Click + drag	Click + drag
Adjust an EQ band's width	Shift + click + drag	Right click + drag -or- Shift + click + drag
Add a new EQ band	Double-click	Double-click
Delete an EQ band	Ctrl + click	Ctrl + click

Show popup menu of all EQ bands listing all parameters	Right-click on display -or- Shift-click on display	Shift-click on display
--	--	------------------------

EQ Tabs

Just below the frequency response display are tabs for selecting what appears in the display, either the Equalizer or the DeNoiser. Click on the Equalizer tab to select the EQ display.

Add and Delete

The "Add" and "Delete" buttons are used to add and delete EQ bands. Clicking Add will create a new EQ band, and will display a new control handle. You can also add a band by double-clicking in the frequency response display. Up to 10 bands may be created. Clicking Delete deletes the current band. You can also delete a band by ctrl-clicking on its handle.

Enable

The Enable button enables/disables the entire EQ section. When disabled, the Enable button is not lit, and the EQ section is bypassed.

Bypass

The Bypass button bypasses the current band only. Toggling the bypass useful to hear the effect of a single EQ band. Any number of EQ bands can be bypassed. As you switch between different bands the bypass button will change to reflect the bypass state of the current band.

Pre and Post

The Pre and Post buttons control whether the EQ operates before the compressors (Pre button lit) or after the compressors (Post button lit). See the routing diagram at the start of this chapter for more information.

EQ Preset

The EQ preset selector button allow you to select factory or user presets for just the EQ section of Dialog. Clicking on the button will step to the next preset. Clicking on the preset name will display a popup menu of the current factory and user EQ presets.

EQ Type

3. Menu Bar and Preset Manager

The EQ Type selector button selects the EQ type of the current band. There are eleven different types, described in the “Equalization” section earlier in this chapter.

Frequency, Height, and Width knobs

These knobs let you set the frequency, height and width of the current band. If the selected EQ type does not support all three parameters, the unused parameters will display “n.a.” (not applicable). For the vintage and resonant shelf filters, the width knob serves as a resonance control by displaying values in units of “Q” rather than octaves.

Note that the EQ Type, Frequency, Height, Width, and Bypass controls will change as you select different bands to reflect the state of the currently selected band. Also note that the EQ preset affects neither the brickwall settings nor the sidechain EQ settings nor the EQ Enable button.

EQ display range

The vertical range of the EQ display can be changed by the user, allowing more subtle editing/viewing of the EQ response. Simply SHIFT-click (or right-click for PC users) on the vertical axis at the left of the EQ display to see a popup menu listing the ranges. The ranges are -24 to 24 dB, -12 to 12 dB, and -6 to 6 dB. Dialog will remember the last used setting. When using the zoomed-in ranges (+/- 12 and +/- 6 dB) it is possible that the EQ control handles will be off the display and will not be drawn. In this case you can select one of these bands by using the SHIFT-click popup menu, and then edit the values using the EQ parameter knobs.

8. Compressor

8.1 About Compressors

Compressors are used to even out dynamics in an audio signal by making loud portions softer and making soft portions louder. Very simply, a compressor tracks the level of the input signal and turns down the gain when the input exceeds a threshold level. Compressors are very useful for dialog processing because dialog recordings often contain too much dynamic variation. There are a number of different scenarios in dialog recording where a compressor can be applied beneficially:

- A single talker can go from shouting to whispering, or may occasionally be too close to the microphone.
- There may be multiple talkers at different distances from the microphone, so each is recorded at a different volume.
- You may have a recording spliced from different sources, say a news clip or interview, where the individual segments are at different volumes.

Threshold and Ratio

A compressor has five principal controls: threshold, ratio, attack time, release time, and makeup gain. The threshold is the input level at which compression will kick in. The ratio parameter determines how much gain reduction will be applied as the input exceeds the threshold. The definition is somewhat archaic and confusing. The range of ratios is 1 to infinity. A ratio of 1 means no gain reduction, while a ratio of infinity means the gain reduction is equal to the amount the input exceeds the threshold (hence the input will be pinned at the threshold value). In general, for a ratio R , the gain reduction is $(R-1)/R$ of the amount the input exceeds the threshold. So to take a typical example, if the input is 12 dB over threshold with a ratio of 3, the gain reduction will be 8 dB.

But compressors are not really about decreasing gain, they are really about increasing gain. What they do is decrease the peaks in a signal so the rest of the signal can be boosted. This is what the makeup gain control does, it applies a constant gain boost to increase the soft parts of a signal while the compressor dynamically pushes down the peaks.

Attack and Release Time

The attack time and release time parameters are really important. The attack time controls how fast the gain is turned down when gain reduction is to be

applied and the release time controls how fast the gain is turned back up. Consider the example of compressing a drum hit. With a small (fast) attack time, the initial transient of the drum hit will be compressed because as soon as the transient exceeds the threshold the gain reduction kicks in immediately. As you increase the attack time, it takes longer for the gain reduction to fade in and the initial transient of the drum hit passes through uncompressed. So in this case, when you want to hear more drum attack, you increase the attack time. After the transient has passed and the drum sound begins its decay, the compressor will begin releasing, increasing the gain back to the nominal level as set by the makeup gain. A short release will restore gain immediately, in which case the drum sound will have essentially the same decay it had originally. A longer release time will restore gain more slowly, thus increasing the decay time of the original drum sound. So the original "BUMPH" of the drum sound becomes "BOOOOMMMMPH" after compressing. That's punch. The ability to change the decay time of acoustic instruments after they have been recorded is a really powerful capability of compressors.

Applied to dialog recordings, compressors are often used as an automatic gain control (AGC). The idea is not to change the character of the voices but to make all the voices approximately the same loudness. Typically a fast attack time (e.g., a few milliseconds) is combined with a relatively slow release time (e.g., half a second). With an infinite ratio, peaks exceeding the threshold cause the gain to be rapidly decreased, and so the compressor is acting like a peak limiter. In the absence of peaks, the slow release time causes a slow gain increase to make quiet sections louder. The threshold and makeup gain determine the range of the AGC operation, e.g. with a threshold of -12 dB and a makeup gain of +12 dB, quiet sections will be boosted by a maximum of 12 dB to match the louder sections.

8.2 Using the Compressor

The user interface for the compressor is shown in the figure below.



Figure 8-1. Dialog Compressor User interface


Clicking ENABLE enables the compressor. Selecting a factory preset is a good place to start. Clicking on the preset selector button will cycle through the factory (and any user) presets. There are factory presets for basic dialog compression (Basic, Gentle, and Hard Compress), gain control (Gain Control and Hard Gain Control), and vocal compression (various "Vox" presets). Clicking on the preset name will show a popup menu of preset options.

The meters are below the preset control. On the left is an input meter showing the level of the input signal. A threshold slider control is displayed next to the input level meter. The peak input level to date is shown at the top of the meter; clicking the value will reset it. On the right is the attenuation meter showing the attenuation of the compressor. The peak attenuation is shown at the top of the meter; clicking the value will reset it. The range of the attenuation meter can be changed by right-clicking (SHIFT-clicking on Mac) the meter axis values, which shows a popup menu of meter ranges.

The controls for makeup gain, threshold, ratio, attack and release times are arranged around the level meters. The function of these controls was described earlier.

9. Output and Limiter

The Dialog output section includes output gain, output meters and peak indicators, and a peak limiter. The user-interface is shown in the figure below.



The left channel meter level is shown on the top and the right channel is shown on the bottom. To the right of the meters are peak indicators that show the peak level to date. These are displayed in red if they are over 0 dB. Clicking on the indicators resets them.

The Dialog peak limiter, when enabled, prevents any peak from exceeding -0.1 dB. It uses a 2 msec lookahead buffer to detect and respond to peaks before they occur, and hence incurs a 2 msec latency when enabled. When it detects a peak that will clip, it turns down the gain very quickly, then slowly restores the gain after the peak has passed. The Dialog limiter is very transparent sounding when used for light limiting, and it can be used to compress heavily. When the peak limiter is enabled, if peaks occur the peak levels are shown in red to the right of the peak indicators. Clicking on the indicators resets them.

10. Installation and Registration

10.1 Installation

Installers for Dialog are found on the downloads page of the Wave Arts website. There are separate installers for Mac and Windows. Mac installers are “.dmg” files which after downloading will expand into a “.pkg” installer file; double-click on the “.pkg” file to launch the installer. Windows installers are “.exe” files; double-click on the “.exe” file to launch the installer. The installers provide various options for selecting which plug-in formats to install and whether to use Pace/iLok or Wave Arts licensing.

10.2 Registration

We support two licensing methods – Wave Arts licensing and PACE/iLok. When installing the plug-ins you must select which version of the plug-ins you wish to use. When you purchase a plug-in, you will be e-mailed a Dialog 2 serial number (looks like WA-DLG2-XXXX-XXXX where DLG2 indicates the serial is for Dialog 2, and X is a hex digit). Use the serial number to unlock the plug-in as described below.

10.3 Wave Arts licensing

Prior to registration, the plug-in operates in demonstration mode; it is fully functional but stops operating after 30 days. To unlock the plug-in after purchasing, open the plug-in, select the Tools->Register option, and enter your name, email address, and serial number. The plug-in will contact our registration server and download a license file which will unlock the plug-in. You should see a message saying your registration was successful.

If your computer is not connected to the internet, use the Tools->Offline Register option and follow the instructions to generate a keyfile at our website registration page and import the keyfile.

10.4 PACE/iLok licensing

All our plug-ins support PACE/iLok. Prior to activation, the plug-ins will allow you to start a 30-day trial by creating an iLok account. To unlock the plug-in after purchasing, go to our Product Registration page, select “PACE/iLok”, and enter your serial number. A PACE redeem code (looks like XXXX-XXXX-XXXX-XXXX-XXXX-XX) will be displayed and also emailed to you. There are two ways to redeem the code and generate a license. When opening the plug-in a dialog window will appear and give you the option to

3. Menu Bar and Preset Manager

Activate the plug-in, you can paste the PACE redeem code there, and proceed to create or login to an iLok account and then transfer the license to an iLok or your machine. Otherwise, go to <http://www.ilok.com>, create an iLok account, and download and install the iLok License Manager. Within the manager, under the Licenses menu, select "Redeem Activation Code" and paste your redeem code. Then transfer the license to either an iLok dongle or your machine. The plug-in will run only if it can find a license on an iLok or the machine.

11. Plug-in Control Operation

This section describes common user interface controls found in Wave Arts plug-ins.

11.1 Knobs



Please refer to the following guide for information about the various ways you can use knobs:

Function	Mac	Windows
Increase/Decrease a parameter value (rotate clockwise/counterclockwise)	Click on the knob + drag up/down*	Click on the knob + drag up/down*
Fine adjustment — increase/decrease	Shift + click + drag up/down -or- Command + click + drag up/down	Right click + drag up/down -or- Shift + click + drag up/down -or- Control + click + drag up/down
Reset knob to default value	Double-click	Double-click

*The up/down drag mode of the knobs can be changed in the Preferences option found in the Tools menu.

11.2 Text Entry



Many value displays are editable text. A text field is editable if your mouse cursor changes to an I-beam when moved over the text. Following is a table that fully describes how to use the text editing features:

Function	Mac	Windows
Enter text entry mode	Click in the display	Click in the display
Select text	Click + drag	Click + drag

3. Menu Bar and Preset Manager

Select entire text	Double-click	Double-click
Delete character to left of cursor	Delete	Backspace
Delete character to right of cursor	Fn+Delete	Delete
Move the cursor left/right	Left/Right arrow keys	Left/Right arrow keys
Extend the current selection	Shift + click + drag -or- Shift + left/right arrow keys	Shift + click + drag -or- Shift + left/right arrow keys
Exit text entry mode	ESC -or- Return/Enter	ESC -or- Return/Enter
Select next parameter to edit	Tab	Tab
Select previous parameter to edit	Shift + Tab	Shift + Tab

You'll find that many parameters, such as frequency, will recognize units typed into the text field. The following values, when typed into a frequency value box, are equivalent:

$$2k = 2 \text{ kHz} = 2000 = 2000 \text{ Hz}$$

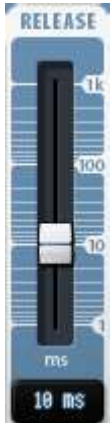
11.3 Selector button





The selector button cycles through a number of fixed values. Click on the button to go to the next value. Click on the text to display a pop-up menu of the available values. The table below describes the functionality of the selector button:

Function	Mac	Windows
Go to next value	Click on the knob	Click on the knob
Go to previous value	Shift + click on knob	Shift + click on knob
Display pop-up menu of all choices	Click on text	Click on text


11.4 Sliders

	Function	Mac	Windows
	Increase/Decrease a parameter value	Click on the slider handle + drag up/down	Click on the slider handle + drag up/down
	Fine adjustment — increase/decrease	Shift + click + drag up/down	Right click + drag up/down -or- Shift + click + drag up/down
	Reset slider to default value	Command + click -or- Double-click	Control + click -or- Double-click

11.5 Buttons

	Lighted buttons show a toggle state. A green, orange or yellow light indicates "on" and a black (extinguished) light indicates "off." Click the button to toggle the state.
	Buttons that do not light up are used to activate certain commands.

11.6 Output Meters

	Output meters show the peak signal power in short time updates — in green from -96dB to -6dB, in yellow between -6dB and 0dB, and in red above 0dB. Peak hold levels are also drawn. The meter also stores the overall peak value for each channel, and displays these values in the peak indicator boxes above the meter. If the detector finds a peak value above 0dB, the text color will turn red as a warning. Click on either indicator box to reset them back to -96dB.
---	--

12. Menu Bar and Preset Manager

All Wave Arts plug-ins in the Power Suite Bundle have the following menu bar displayed at the top of the plug-in:



This section describes the operation of the menu bar, preset manager, and the other functions available in the menus.

12.1 Bypass

Clicking on the bypass button bypasses the effect, that is, audio will pass through the effect without alteration. The button is lit when the effect is bypassed.

12.2 Undo

Clicking the Undo button causes the parameters to revert to their settings prior to the last edit. Only one level of undo is available, so clicking the undo button again will restore the parameters after the edit. Both A and B buffers (described below) have their own undo buffers.

12.3 Copy

Clicking the Copy button copies the current set of effect parameters to the unused A/B buffer. Hence, if the A buffer is currently selected, the parameters are copied to B, and if the B buffer is selected, the parameters are copied to A. After clicking Copy, you can continue to make changes, and then revert to the original copied settings by clicking either the A or B buttons to switch buffers.

12.4 A/B buffers

The A/B edit buffers allow you to compare two different sets of parameters or presets. One of the A or B buttons is always lit; the button that is lit shows the current buffer. Clicking either the A or B button will switch to using the other buffer, thus changing the effect settings (assuming different settings are stored in A and B).

Here's how to use the A/B buffers to compare two different presets. Select a preset from the Preset menu, then switch to the other buffer and select a different preset. Now switch between the two buffers to alternate between the two different presets.

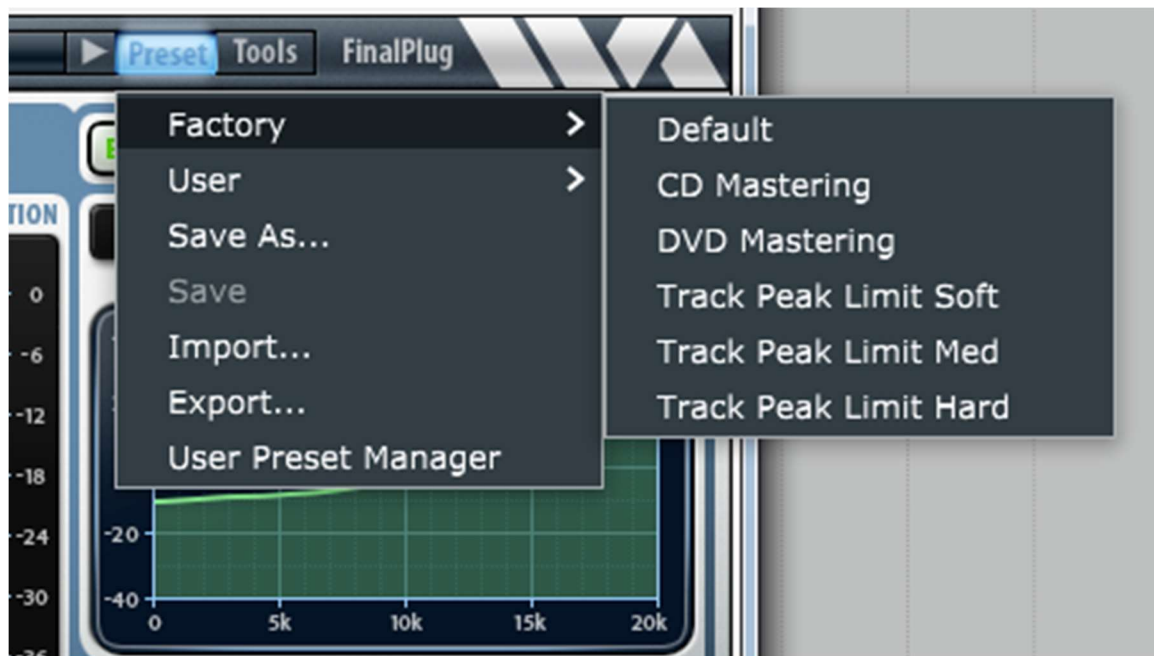
12.5 Preset name and arrow controls

The currently selected preset name is displayed in the text field in the menu bar. Changing any parameters causes an asterisk (*) to be displayed at the end of the name. This indicates that changes have been made to the preset. In order to save the changes to a user preset you must select the "Save..." item in the Preset menu, described below.

The arrow controls to the left and right of the preset name cycle through the set of factory and user presets. Clicking the right arrow goes to the next preset, clicking the left arrow goes to the previous preset.

12.6 Preset menu

The Preset menu contains lists of factory and user presets for easy selection, and options for managing presets. The functions are described in the following sections.



12.7 Factory Presets

Factory presets are selected from a rolloff menu at the top of the Preset menu. Factory presets cannot be modified or deleted. The Default preset is always first in the list; it defines all default parameter settings.

12.8 User Presets

User presets are selected from a rolloff menu just below the Factory presets in the Preset menu. When you first run a Wave Arts plug-in, there will not be any user presets and the menu will be empty. When you save a preset using the "Save" option the preset is added to the User menu. All instances of a plug-in share the same set of user presets. So, after you save a preset with one instance of a plug-in, you can go to another instance and find that the preset can be found in its User preset menu too.

User presets are stored in a text file called "<plugin> Presets.txt", where <plugin> is the name of the plug-in you are using. If the file is deleted, an empty preset file will be created automatically the next time the plug-in runs. User presets files are stored in the following directory, depending on the operating system, where <username> is your Windows login name:

Mac OS-X:

/home/<username>/Library/Application Support/Wave Arts/<plugin>/

Windows:

C:/Users/<username>/AppData/Roaming/Wave Arts/<plugin>/

12.9 Save As...

When you have created an effect you want to save as a preset, select the "Save As..." option. You will be asked to name the preset and the preset will be saved in the set of User presets. If you supply the same name as an existing user preset, the preset will be overwritten with the new preset without any warning notice.

12.10 Save

The Save option is used to save changes to a user preset. The Save option is enabled when the current preset is a user preset and you have made changes, in which case an asterisk (*) will be appended to the preset name. Select the Save option to save the changes to the preset.

12.11 Import...

User presets can be written to files using the "Export" function, and read from files using the "Import" function. Selecting the "Import..." option will open a file chooser to select the preset file for importing. After selecting the file the presets are read and will appear in the User Presets menu. Importing does not delete existing presets; if an imported preset has the same name as an existing preset, you will get two presets with the same name.


12.12 Export...

Selecting the “Export...” option will open a file chooser to specify the output file. Then your user presets are written to the file.

Preset Export is also useful for making backup copies of your user presets. If you have a large set of user presets, be sure to export them to a backup file.

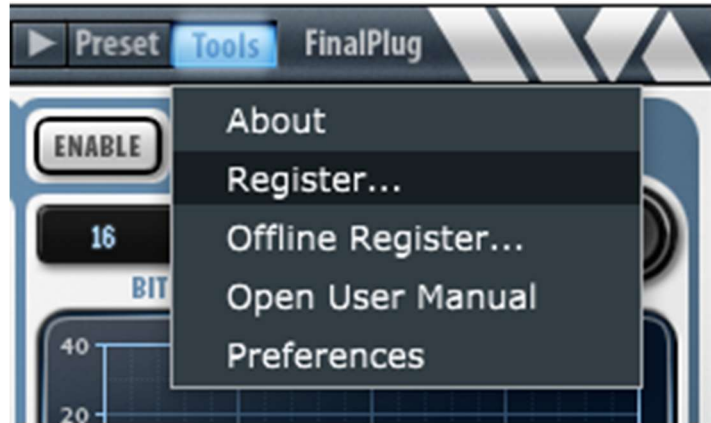
12.13 User Preset Manager

The User Preset Manager provides a simple interface to organize user presets. You can rename presets, delete presets, and organize them by moving them up/down in the list.




12.14 Tools menu

The Tools menu contains various important options, described below.



12.15 About...


The About option displays important information about your plug-in. An example is shown below:



On the top line, the plug-in name and version are displayed, along with the current plug-in format (AAX, VST3, AU). This is useful if you aren't sure which format of the plug you are running. The build date of the plug-in is displayed on the next line. If the plug-in is using Wave Arts licensing, the registration status is displayed on the next line. If the plug-in is operating in demo mode, the time remaining (if any) is displayed. If the plug-in has been successfully registered (unlocked), the serial number is displayed. If the plug-in is using Pace/iLok licensing, it will display "Pace/iLok licensing".

12.16 Register...

Select this option to register (unlock) your plug-in. Enter your name, email address, and serial number, and click OK.

A dark-themed dialog box titled "Register" with a red close button in the top right corner. The dialog contains the instruction "To register, enter your serial number, name, and email address." followed by three input fields labeled "Serial #:", "Name:", and "Email:". At the bottom, there are two buttons: "Cancel" and "OK".

Register

To register, enter your serial number, name, and email address.

Serial #:

Name:

Email:

Cancel OK

12.17 Offline Register...

Select this option to register the plug-in when your computer is not connected to the internet.

12.18 Open User Manual...

Select this option to open this user manual in a browser.

12.19 Preferences...

This option opens the Preferences dialog to customize the plug-in operation.

License Agreement

END-USER LICENSE AGREEMENT FOR WAVE ARTS SOFTWARE

IMPORTANT-READ CAREFULLY: This Wave Arts End-User License Agreement ("EULA") is a legal agreement between you (either an individual or a single entity) and Wave Arts, Inc. ("Wave Arts") for the Wave Arts software accompanying this EULA, which includes computer software and may include associated media, printed materials, and "on-line" or electronic documentation ("SOFTWARE PRODUCT" or "SOFTWARE"). By exercising your rights to make and use copies of the SOFTWARE PRODUCT, you agree to be bound by the terms of this EULA. If you do not agree to the terms of this EULA, you may not use the SOFTWARE PRODUCT.

SOFTWARE PRODUCT LICENSE

The SOFTWARE PRODUCT is protected by copyright laws and international copyright treaties, as well as other intellectual property laws and treaties. The SOFTWARE PRODUCT is licensed, not sold.

GRANT OF LICENSE This EULA grants you the following rights:

Installation and Use.

When using Wave Arts licensing, you may install and use the SOFTWARE PRODUCT on up to three computers, provided you are the only user. A license must be acquired for each individual user of the SOFTWARE PRODUCT. When using Pace/iLok licensing you are granted a single user license which you can transfer to an iLok dongle or to a machine.

Reproduction and Distribution.

You may not reproduce or distribute the SOFTWARE PRODUCT except to make backup copies, or to install as provided for above.

DESCRIPTION OF OTHER RIGHTS AND LIMITATIONS

Limitations on Reverse Engineering, Decompilation, and Disassembly.

You may not reverse engineer, decompile, or disassemble the SOFTWARE PRODUCT, except and only to the extent that such activity is expressly permitted by applicable law notwithstanding this limitation.

Software Transfer.

You may permanently transfer all of your rights under this EULA, provided you retain no copies, you transfer all of the SOFTWARE PRODUCT, and the recipient agrees to the terms of this EULA.

Termination.

Without prejudice to any other rights, Wave Arts may terminate this EULA if you fail to comply with the terms and conditions of this EULA. In such event, you must destroy all copies of the SOFTWARE PRODUCT and all of its component parts.

COPYRIGHT

All title and copyrights in and to the SOFTWARE PRODUCT (including but not limited to any images, text, and "applets" incorporated into the SOFTWARE PRODUCT), the accompanying printed materials, and any copies of the SOFTWARE PRODUCT are owned by Wave Arts or its suppliers. The SOFTWARE PRODUCT is protected by copyright laws and international treaty provisions. Therefore, you must treat the SOFTWARE PRODUCT like any other copyrighted material.

LIMITED WARRANTY

NO WARRANTIES. Wave Arts expressly disclaims any warranty for the SOFTWARE PRODUCT. The SOFTWARE PRODUCT and any related documentation is provided "as is" without warranty of any kind, either express or implied, including, without limitation, the implied warranties or merchantability, fitness for a particular purpose, or noninfringement. The entire risk arising out of use or performance of the SOFTWARE PRODUCT remains with you.

NO LIABILITY FOR CONSEQUENTIAL DAMAGES. In no event shall Wave Arts or its suppliers be liable for any damages whatsoever (including, without limitation, damages for loss of business profits, business interruption, loss of business information, or any other pecuniary loss) arising out of the use of or inability to use this Wave Arts product, even if Wave Arts has been advised of the possibility of such damages. Because some states/jurisdictions do not allow the exclusion or limitation of liability for consequential or incidental damages, the above limitation may not apply to you.

MISCELLANEOUS

If you acquired this product in the United States, this EULA is governed by the laws of the State of Massachusetts.

If this product was acquired outside the United States, then local laws may apply.

Support

For assistance, please send email to:

support@wavearts.com

If you are having problems with a plug-in, please include the following information: plug-in name, operating system, and host software you are using. The version numbers are also helpful.

For software updates, revision history, frequently asked questions (FAQ), and more, please visit our website at:

www.wavearts.com

Trademark Notices



VST is a trademark of Steinberg Media Technologies GmbH

Index

A

A/B buffers.....	39
About	43
Acoustic noise.....	12
AGC.....	31
Amount.....	15
Attack time	30

B

Bandpass filter.....	25
Brickwall filter.....	8
Broadband noise	12
Buttons.....	38
Buzz.....	10
Bypass.....	39

C

Compressor.....	30
Copy.....	39

D

De-Ess.....	19
De-Hum	10
De-Noise	12
De-Plode	16

E

EQ.....	22
---------	----

F

Factory presets	40
Floor.....	12

H

Harmonics.....	10
Highpass filter	8, 26
Hum	10

I

iLok.....	34
Import preset	41
Installation	34

K

Knobs	36
-------------	----

L

Learn.....	13
------------	----

Licensing.....	34
Limiter.....	33
Lowpass filter	8, 26

M

Magnetic tape	12
Makeup gain	30
Menu bar	39
Meters	38
Monitor	18, 21

N

Noise	12
Notch filter	25

O

Output.....	33
-------------	----

P

PACE	34
Parametric EQ.....	22
Peak limiter	33
Plosives	16
Pre/Post.....	6
Preset menu	40
Preset name	40
Presets	40

R

Ratio	30
Registration.....	34
Release time.....	30

S

Save preset.....	41
Selector button	37
Shelf filter	23
Sibilants.....	19
Sliders	38
Spectrum	12, 14

T

Tabs	28
Text entry	36
Threshold.....	30
Tools menu	42

U

Undo	39
------------	----

User manual 44

User presets 41