

Waves Z-Noise



Software Audio Processor Users Guide

1. Introduction

Waves Z-Noise is a single-ended broadband noise reduction audio processor. It very effectively reduces noise while preserving the highest audio quality. Superficially, Z-Noise works like other broadband noise reduction processors which use a noise profile, or noise print, to differentiate valid signal from noise. However, Z-Noise uses intelligent multi-level decision algorithms to maximize noise reduction with minimum unwanted artifacts. Plus, Z-Noise uses a more musical band split method for more natural results throughout the sound spectrum.

Z-Noise is ideal for removing additive noise such as tape hiss, vinyl surface damage, wind, air conditioner and ventilation system noise. Z-Noise offers several improvements, including:

- Z-Noise works on the complete frequency spectrum, thus high-frequency hiss and low-frequency rumble can be removed in the same pass. It is a truly *broadband* noise removal tool.
- *Real-time* operation allows you to adjust parameters and hear the results immediately.

There are several new features in Z-Noise which address problems unique to broadband noise reduction:

- **Extract Profile Mode** – Not always do you have a section of pure noise from which to create a noise profile. Using the Extract Noise Profile feature, you can create a noise profile from a noise sample which is “contaminated” by program material which you wish to save.
- **Adaptive Mode** - You do not necessarily need to create a noise profile in order to reduce noise. In the Adaptive Mode, Z-Noise extracts a noise profile on-the-fly and reduces the noise even if the profile changes over time.
- **Better transient treatment** – Z-Noise offers improved transient handling as well as enhanced transient processing options that allow for greater noise reduction with less transients smearing. Z-Noise can detect and enhance transients prior to noise reduction processing, then de-enhance the transients after noise reduction, thus allowing for more aggressive noise reduction while preserving transients.
- **Enhanced Noise profile adjustments** – In both the normal and adaptive modes, Z-Noise allows you to control the spectral nature of the noise profile. In this manner you can more precisely define what is a noise and what is program material. The noise profile can be equalized using a five-band parametric EQ interface.

- **Knee control** – Here you adjust the angle of the knee to smooth or sharpen reduction dynamics.

How Does Z-Noise Work?

Z-Noise, like other broadband noise reduction algorithms, is essentially a multi-band expander that uses a noise profile to set the threshold within each band according to the energy of noise analyzed in that band. When energy in a specific band falls below the threshold, the input signal is presumed to be noise or signal decaying into noise and the signal is attenuated in that frequency band. Z-Noise employs the most recent psychoacoustic research and multilevel decision algorithms to remove noise while preserving the sonic clarity of the source. It performs much better than comparable DAW tools and as well as or better than more expensive hardware solutions. Used wisely, Z-Noise reduces noise while minimizing artifacts, and preserves the quality of the original recording.

In the case of broadband noise reduction, noises are considered to be uncorrelated from the musical signal. Z-Noise will do a great job of removing tape flutter or analog hiss, for example, but it will not be able to reduce background voices or anything that might be interpreted as signal, such as birds, crickets, music, etc.

The Z-Noise is a single-ended noise reduction algorithm that does not require an encoded original recording. Instead, it learns or extracts a *noise profile* from the audio signal and uses this noise profile to help it to distinguish noise from audio data.

Z-Noise's controls are similar to those of a dynamics processor. Those familiar with a typical compressor/expander should achieve excellent results within a few minutes of experimentation.

2. Using Z-Noise

De-noising with Z-Noise is a two-step process. First, you create a valid noise profile, or signature, of the noise afflicting your audio. Then, use Z-Noise's many de-noising tools to obtain the perfect noise reduction to suit your program material.

Step One – Create a Noise Profile

You must obtain a noise profile in order to effectively reduce noise. A noise profile is the spectral density estimation of the noise, and is used by Z-Noise's noise suppression algorithms to make complex decisions as to what is a noise, and should be eliminated, and what is not a noise, and should be left untouched.

Z-Noise provides five methods for creating a noise profile, depending the nature of your original material and your desired work methods.

Default profile

When you start Z-Noise, a default white noise profile is loaded. This white noise profile also loads when the *Full Reset* preset is chosen.


Factory presets

Other standard noise profiles may be loaded from factory presets.

Learn


Use the Learn mode when you have a clean sample of pure noise, with no program material intruding upon the noise. Those familiar with Waves X-Noise will be comfortable with this process.

Select a section of audio, at least 100ms long, which contains only the noise you are trying to eliminate. Be careful to avoid "contaminations" from your source audio, such as ring outs and false attacks.

Click the *Learn* button in the Noise Profile area below the Z-Noise analyzer. The button will blink "Learning" . Play the selected audio section through Z-Noise. Click *Learn* again to stop the learning process and create the noise profile, which appears as a white line on the Z-Noise analyzer. This line represents the frequency content of the analyzed noise and changes its vertical position as you modify the *Threshold*. At this point, save your Z-Noise setup, which includes your noise profile. (For more on saving and recalling set-ups, refer to the **WaveSystem** section at the end of this manual.)

If you cannot find a noise-only section, try using the default white noise profile or one of the other available factory presets, or use the Extract Profile process described below.

Extract

This option will enable you to create a fixed noise profile out of an audio section where no “pure noise” is available, as you might encounter when cleaning music on vinyl disks. To use this option, select an audio section in which the signal is relatively quiet and whose program material “contamination” is representative of the rest of program. Then click the *Extract* button. The button will blink  “Extracting.” Play the selection of audio. When finished, click the button again to see the extracted noise profile.

Since Z-Noise will need to create an average calculation of the non-noise signal in order to construct a valid noise profile, you should work with a relatively long and representative sample when using the Extract Profile mode. Remember, both the Learn and Extract methods are designed for recordings with stable noise backgrounds. For modulating or changing background noises you should choose the Adaptive method of creating noise profiles.

Adaptive



The Adaptive mode is used to eliminate noise that changes over time. Z-Noise updates the noise profile by analyzing the input signal on-the-fly. Click the *Adaptive* button, then play your DAW to sample the sound you want to clean. Z-Noise will start scanning the input audio and automatically assign a noise profile (the initial noise profile will be a white noise profile or the last profile captured.) The noise profile appears as a white line in the Analyzer window.

The Adaptive mode requires about two seconds of audio throughput before it begins creating its own noise profile. During this initialization period, noise reduction can be less than ideal, since Z-Noise lacks a valid noise profile. In order to get optimal results from the beginning of your sound source, experiment with these techniques:

- 1) Use the **Learn** mode to create a noise profile of a section of pure noise at or near the beginning of your source audio. Z-Noise will begin Adaptive mode noise reduction with this Learn mode noise profile, then migrate to the Adaptive profile in a smooth way. This way you will have effective noise reduction from the beginning of your source material, plus the constantly updating noise reduction offered with the Adaptive mode.
- 2) If no pure noise is available at the beginning of your source audio, you can use the **Extract Profile** mode to create a noise profile from a segment containing both musical signal and noise. As in the previous example, this noise profile will become the starting point in the Adaptive noise reduction process.
- 3) If you don't create a noise profile prior to starting an Adaptive mode noise reduction process, Z-Noise will default to a white noise profile until it has accurately created an Adaptive noise profile. After about two seconds, it will shift to its new noise profile.

To reset the noise profile, go to the **load** menu and load the *Reset* preset (either *Both* or just *Profile*.)

It is important to create a noise profile and perform noise reduction at the same sample rate. Doing otherwise can result in less accurate noise reduction.

Step Two – Noise Reduction

Once you have created a noise profile of your audio, you are ready to remove the noise. Z-Noise uses two primary parameters, *Threshold* and *Reduction*, to control the noise reduction process. Adjust these controls first to quickly configure Z-Noise for a given input signal. *Attack*, *Release*, *Transients*, *Knee*, *Optimize* and *Profile EQ* parameters offer more detailed control and are described later.

Threshold - sets a global offset for the noise profile. Here you set a level used to begin distinguishing between the source sound and the noise. Setting the *Threshold* to a high setting, such as +10 dB, means that the noise profile is offset upwards by 10dB, resulting in more signal subject to removal.

Reduction - Use the *Reduction* control to set the amount of noise reduction applied to signals falling below the threshold. Increasing the *Reduction* setting increases the amount of noise removed from below the noise profile.

If time-aliasing artifacts (singing or robot-like sounds) appear, decrease the *Reduction* setting or lower the *Threshold* to several dB above the background noise. Artifacts are often an indication that program material is being processed along with noise. Lowering the Threshold setting “protects” the program material from unnecessary processing.

Artifacts can be further minimized by adjusting the *Attack*, *Release*, *Optimize*, *Knee*, *Transients* and *EQ* parameters (see the **Controls** section for more details.)

For a complete user interface reference, see **Controls and Displays** later in this manual.

Step Three – Monitoring

A useful Z-Noise feature is the ability to monitor either the de-noised *Audio* output or the *Difference*, the noise signal being removed by Z-Noise. Listen carefully to the *Difference* signal to determine whether your settings are removing parts of the original audio signal in addition to the noise. If you hear elements of the original signal in the *Difference* monitor mode, you are likely degrading the

audio signal.

The goal is to maximize noise reduction while affecting as little of the original program material as possible. In addition to the *Threshold* and *Reduction* controls, there are several tools you can use to find the best set of parameters for your program material. These controls will be discussed in the next section.

Waves recommends monitoring with headphones and alternating between *Audio* and *Difference* monitor modes several times until the right set of parameters is found. Heavily damaged recordings may require compromise between audio quality and noise reduction.

3. Controls and Displays

Controls



The two fundamental noise reduction controls are *Threshold* and *Reduction*.

Threshold - Settings: -20 to +50 dB; Default = 0 dB

This is the reference gain setting, as compared to the noise profile. Signal below the noise profile is cleaned, while the signal above the profile is ignored or minimally processed.

Reduction - Settings: 0 - 100%; Default = 0 %

Determines the amount of noise reduction applied to the signal below the *Threshold*. The higher the setting, the greater the noise reduction.

Dynamics Controls



Attack - Settings: 0.01–1000ms; Default = 0.01ms

Attack is the time it takes for attenuation to stop or for the gain to rise. Noise reduction is smoothly decreased to avoid pops and clicks from abrupt processing changes. The default setting (0.03 second) works well in most cases. Explosive sounds may require a shorter time; sounds that evolve more slowly may require a longer *Attack* time.

Release - Settings: 0.01 - 5000ms; Default = 0.40 ms

Sets the time it will take for noise reduction or gain attenuation to occur or to reach its target reduction value. Higher *Release* values will result in slower attenuation; these settings will tend to sound smoother but will likely let more noise through than will short values. Short *Release* values are usually more effective in removing noise, but may result in an edgier sound.

Transients - Settings: 0 dB (= off) to +15 dB Default = 0.1dB

Transient control is a unique feature of Z-Noise which preserves transients by enhancing them before the noise reduction process takes place. This enables you to increase the *Threshold* without compromising transients. When noise reduction is completed, the transients are de-enhanced. Note that you may need to re-adjust the *Threshold* control when increasing the *Transient* settings.

Knee - Settings: 0=Hard, 100=Soft, Default 25%

The *Knee* control adjusts the sharpness of the gain reduction curve below the noise profile. With a hard knee setting, the dynamics will activate when energy

passes the threshold. With softer knee, dynamics will begin to activate as energy approaches the threshold.

Noise Profile EQ Controls

The *Noise Profile EQ* section lets you manipulate the noise profile using a five-band parametric equalizer. This allows you to control the amount of noise reduction in different frequency ranges. Its effect is similar to that of the *Threshold* slider, except that it is frequency dependent.

Noise Profile EQ Parameters:



The *Profile EQ* allows you to change the frequency characteristics of your noise profile. The five bands of EQ control, as indicated by five numbered square markers on the noise profile plot, are always visible and accessible from the Z-Noise graph in the Analyzer window. When you select the marker of a noise profile EQ band, its parameters will appear in the EQ controls section just below the graph. This allows for parametric control. When selecting multiple bands, the first band marker selected (and active) will be the one that appears in the profile EQ strip below the graph.

The Graph markers are numbered and color coded:

Band 1	red	default frequency 60Hz	default type is low shelf.
Band 2	magenta	default frequency 181Hz	default type is bell.
Band 3	blue	default frequency 577Hz	default type is bell.
Band 4	green	default frequency 1702Hz	default type is bell.
Band 5	yellow	default frequency 5806Hz	default type is high shelf.

Each band offers the following controls –

Band On/Off – Default = On

This determines whether the Profile EQ band is engaged or bypassed. Dragging the band marker on the graph engages the frequency band. A graph marker will highlight when on and dim when off.

Band Type – low shelf, notch, bell, high shelf.

Default band 1: low shelf; default bands 2-4: bell; default band 5: high shelf

This determines the shape of the filter for each band. The notch filter (a bell with a very high Q) is useful when removing distinct hum harmonics.

Frequency - Settings: 10 Hz - 20 kHz

Sets the frequency in each of the five bands of the Noise Profile EQ. The frequency of the active marker is displayed.

Gain - Settings: -30dB to +30 dB; Default = 0 dB

Gain controls the attenuation or boost applied to the noise profile frequencies. Increasing a frequency's *Gain* raises the threshold for that frequency, causing the algorithm to reduce more noise in that spectrum. Decreasing *Gain* lowers the threshold for that frequency, resulting in less noise reduction in that frequency.


Q - Settings: 0.26 (wide curve) - 6.5 (sharp curve); Default=1.0

Sets the Q of the frequency to be reduced or boosted. A higher Q results in a narrower filter, while a lower Q yields a wider, gentler filter.

Resolution Optimization Controls

Broadband noise reduction necessitates a balance between resolution in the time domain vs. resolution in the frequency domain. There's no perfect solution, as diverse program material and various noise problems require unique compromises.

Z-Noise allows you to choose between higher frequency resolution or higher time resolution, or a compromise between the two, thus optimizing resolution to best suit your program material.



Optimize - Settings: Punch, Norm, Smooth; Default = Norm

Allows you to choose between higher frequency or time resolution:


Smooth sets the highest frequency resolution. In the *Smooth* mode, Z-Noise provides the smoothest and finest overall noise reduction but with impaired transient quality. Also, this mode consumes more CPU power and may not be suitable for slower computers. This mode is best used for voice and other sources without sharp attacks.

Punch sets the best time resolution and is recommended for materials that have significant attacks and transients, such as percussive sounds or plucked strings. Piano recordings are usually best noise reduced when optimized for better *Punch*.

Norm offers a nice compromise for many situations, with better transients than in *Smooth* and better frequency resolution than in *Punch*.

Waves recommends creating a noise profile and processing the sound source with the same resolution optimization, although it is possible to change resolution between these actions.

Noise Profile Creation Controls




The *Noise Profile* is a fingerprint of the noise. It is an essential element of effective noise reduction, as this is what enables Z-Noise to classify audio as noise or signal in a specific recording. Z-Noise offers three ways to profile noise:

Learn; *Extract*; and *Adaptive* modes.

Learn



Best results are obtained when the noise is static and you use a small segment of “pure noise” from your source audio. Use the *Learn* function with a selection consisting of *only* the background noise you want to remove. Click the **Learn** button to begin the process; the button changes to “Learning”  and blinks until it is clicked again. *Learn* assumes the input is pure noise and profiles it with no modifications.

Extract



If your program material doesn't have a section with pure background noise, find a section that is mostly noise and use *Extract* in the same way you used the *Learn* button to create a noise profile. *Extract* assumes the input includes some noise and some signal. It uses advanced intelligent algorithms to extract a noise profile from a section of audio that does not have a pure noise segment from which to “learn” a profile.

It is important to emphasize that no noise reduction occurs while creating the noise profile in the *Learn* or *Extract* modes.

Adaptive



Both the *Learn* and *Extract* modes of noise profile creation assume that your noise is steady-state, i.e., it does not vary in character over time. There are, however, times when Z-Noise must react to an ever-changing noise profile. In these difficult circumstances, Z-Noise's unique, on-the-fly *Adaptive* mode is vital. In this mode, Z-noise will use the most recently learned or extracted noise profile as a starting point in noise reduction. If no profile was learned/extracted, the starting point will be a white noise profile.

Output Monitor Control



The Output Monitor switches between *Audio* (the audio after Z-Noise processing) and *Difference* (the noise being removed with the selected settings.) *Audio* is the default setting. Use the *Difference* monitor setting to hear the noise being removed from the source signal. If you hear program material within the *Difference* signal, modify your settings to achieve a better balance between noise reduction and signal loss/degradation.

Displays

Z-Noise Analyzer



The Z-Noise Analyzer is the processor's main display. It shows four colored spectral envelopes:

- **Red** the input signal before Z-Noise processing
- **White** the noise profile
- **Yellow** the profile shaping EQ curve.
- **Green** the output signal after Z-Noise processing

In a healthy noise reduction process, the green output signal line will generally fall below the red input signal line. However, peaks may overlap the red and green lines.

NR: Noise Reduction Meter

The Noise Reduction Meter shows the energy of noise being removed. The meter indicates what is being removed from the audio signal, and the view is unaffected by the setting of the **Output Monitor Control** switch.

When listening in the *Difference* mode, the Noise Reduction Meter reflects the removed noise you are hearing.

4. Noise Profile Management

What is a Noise Profile?

A Noise Profile is a representation of the noise targeted for removal, shown as a frequency response plot. Z-Noise uses the noise profile to analyze and process input audio to remove noise which has the same characteristics as the noise profile.

Creating a Noise Profile

Z-Noise's *Learn* and *Extract* modes require analysis of a section of static noise in order to create a noise profile. In the Adaptive mode, no learning of a noise profile is required, as this calculation happens on-the-fly.

To create an effective noise profile, locate a section of the source recording (at least 100 ms) containing pure noise. These sections are typically found immediately before the audio begins, after it ends, or during a pause in speech or music. If you take a sample from the end of a recording or during a pause, be careful of ring-out in the recording, as this will compromise the noise profile.

To use the Learn mode, press the *Learn* button then play the noise segment in your DAW. The *Learn* button flashes "Learning" to indicate that Z-Noise is creating the noise profile. When you finish playing the sample, press the *Learn* button again to stop the learning process and to complete the noise profile. The spectrum of the noise profile is displayed in the Z-Noise Analyzer as a white line.

If the source material does not have a section containing the pure noise you want to remove, you can move to Extract Profile mode and create a noise profile. Although this method of noise profile creation is essential when you are unable to locate a pure noise sample, don't make *Extract* your first choice. This method generally compromises noise reduction efficiency because isolating the noise from the audio is less accurate and can lead to audible artifacts.

Important notes: The noise profile should be created from the same source recording as the material you are noise reducing with Z-Noise. Creating a noise profile from a different source will not help Z-Noise recognize the noise in your source. However, if you cannot find a noise-only segment in your source but can find one from the same recording session under the same conditions, you can use it instead and obtain good results.

Note that when cleaning previously released material, such as vinyl albums or audiocassettes, the noise in pauses between songs is not necessarily representative of the noise of the recordings.

At this point, we recommend saving your Z-Noise setup, which includes the noise profile.

Saving, Loading and Sharing Noise Profiles

The **WaveSystem** includes a *Save* button that stores the noise profile in the setup file along with other parameter data. Each Z-Noise setup file has two spaces available for noise profiles, allowing analysis and storage of two segments of noise which can be easily compared to see which works best.

After creating a noise profile, an asterisk indicates that the setup has been modified in the Setup A/B name bar. This noise profile can be saved into a new setup or over the current one. A noise profile can be loaded from any previously saved setup from the Load menu. This can be useful if you need to process other sessions from the same recording conditions.

5. Important information and Useful Tips

Latency

To properly perform its task, Z-Noise must look ahead to into the future. It manages this incredible feat by delaying the source signal by 34,702 samples. Z-Noise latency translates into the following time delays, depending on sample rate:

44.1kHz	0.787 sec.	This latency is particularly significant when the noisy track runs along with other tracks. To maintain sync, other tracks should be delayed by the same amount. To use Z-Noise in a digital audio workstation, it is important that the DAW is able to compensate for the plug-in's delay. If this latency compensation feature
48kHz	0.723 sec.	
88.2kHz	0.393 sec.	
96kHz	0.361 sec.	

is not available in your host DAW, add at least 34,702 samples of silence at the end of the file and trim the start after the process is completed. Some multi-track host applications perform latency compensation automatically.

Since there is no practicable way to delay your other tracks in a live recording, Z-Noise is not recommended when monitoring a live event.

Storing and Recalling the Noise Profile

Some audio host applications cannot store the noise profile data with the other plug-in settings. This means that saving a high level document which includes Z-Noise processing (such as a session, song or other work) may not guarantee proper saving of the associated noise profile. We recommend saving setups within the WaveSystem to insure that the noise profile is safely stored.

Side Effects

Used properly, Z-Noise creates minimal audible artifacts. Artifacts which can occur include time aliasing artifacts (sometimes referred to as *leftovers*, *gremlins*, *singing robots*, or *blips*) that sound like whistles or lingering oscillations. These can be treated by first lengthening the *Attack* and/or *Release* times. Selecting a softer *Knee* setting can reduce artifacts, as can using the *Smooth* optimize setting. If the artifacts remain, choose more moderate *Threshold* and *Reduction* settings.

Host Application Buffer Size

Waves recommends setting the buffer size in your host application to at least 1024 samples. This will prevent CPU "spikes" during moments of especially intense processing.

Very Loud Program Material

Noise removal with Z-Noise can result in very small level increases (fractions of a dB) in short transients. Whereas this will not affect perceived levels, it could cause a clipping problem with program material whose original levels reach 0dBFS. If your source material has transients peaking at 0, we suggest lowering the level by perhaps 1dB BEFORE beginning to process, to allow for these new peaks.

Don't Change Sample Rate

For best results, use the same sample rate for creating a noise profile and for noise reduction. Switching between sample rates between noise profile creation and noise reduction can result in small but unwanted alterations in the noise profile, and may compromise results.